

---

## BACKWARD REACHABILITY OF ARRAY-BASED SYSTEMS BY SMT SOLVING: TERMINATION AND INVARIANT SYNTHESIS

SILVIO GHILARDI<sup>1</sup> AND SILVIO RANISE<sup>2</sup>

<sup>1</sup>Dipartimento di Scienze dell'Informazione, Università degli Studi di Milano (Italy)  
*e-mail address:* ghilardi@dsi.unimi.it

<sup>2</sup>FBK-Irst, Trento (Italy)  
*e-mail address:* ranise@fbk.eu

---

**ABSTRACT.** The safety of infinite state systems can be checked by a backward reachability procedure. For certain classes of systems, it is possible to prove the termination of the procedure and hence conclude the decidability of the safety problem. Although backward reachability is property-directed, it can unnecessarily explore (large) portions of the state space of a system which are not required to verify the safety property under consideration. To avoid this, invariants can be used to dramatically prune the search space. Indeed, the problem is to guess such appropriate invariants.

In this paper, we present a fully declarative and symbolic approach to the mechanization of backward reachability of infinite state systems manipulating arrays by Satisfiability Modulo Theories solving. Theories are used to specify the topology and the data manipulated by the system. We identify sufficient conditions on the theories to ensure the termination of backward reachability and we show the completeness of a method for invariant synthesis (obtained as the dual of backward reachability), again, under suitable hypotheses on the theories. We also present a pragmatic approach to interleave invariant synthesis and backward reachability so that a fix-point for the set of backward reachable states is more easily obtained. Finally, we discuss heuristics that allow us to derive an implementation of the techniques in the model checker MCMT, showing remarkable speed-ups on a significant set of safety problems extracted from a variety of sources.

---

*1998 ACM Subject Classification:* D.2.4, F.3.1, I.2.2.


*Key words and phrases:* Infinite State Model Checking, Satisfiability Modulo Theories, Backward Reachability, Invariant Synthesis.

This paper extends [39] with all the proofs and adapts materials in [37, 41] to make it self-contained.

## CONTENTS

1. Introduction	3
1.1. Array-based systems and symbolic backward reachability	3
1.2. Symbolic backward reachability and Invariant synthesis	5
How to read the paper	6
2. Formal Preliminaries	7
2.1. Case Defined Functions	7
2.2. Embeddings	8
2.3. A many-sorted framework	9
3. Backward Reachability	9
3.1. Array-based Systems	9
3.2. Backward Reachable States	11
3.3. Tableaux-like Implementation of Backward Reachability	13
4. Termination: a semantic analysis	15
4.1. Configurations	16
4.2. Undecidability of the safety problem	16
4.3. Decidability of the safety problem: sufficient conditions	16
5. Invariants Search	21
5.1. Safety Invariants	21
5.2. Invariant Synthesis	24
6. Pragmatics of Invariant Synthesis and Experiments	32
6.1. Integrating Invariant Synthesis within Backward Reachability	32
6.2. Index Abstraction	34
6.3. Signature Abstraction	34
6.4. Experiments	35
7. Discussion	37
7.1. Related work	39
7.2. Future work	41
Acknowledgments	41
References	41
Appendix A. Omitted Proofs	44
Decidability of restricted satisfiability checking	44
Undecidability of backward reachability	45
Undecidability of unrestricted satisfiability checking	47

## HOW TO READ THE PAPER: MAIN TRACKS



## 1. INTRODUCTION

Backward reachability analysis has been widely adopted in model checking of safety properties for infinite state systems (see, e.g., [1]). This verification procedure repeatedly computes pre-images of a set of unsafe states, usually obtained by complementing a safety property that a system should satisfy. Potentially infinite sets of states are represented by constraints so that pre-image computation can be done symbolically. The procedure halts in two cases, either when the current set of (backward) reachable states has a non-empty intersection with the set of initial states—called the *safety check*—and the system is unsafe, or when such a set has reached a fix-point (i.e. further application of the transition does not enlarge the set of reachable states)—called the *fix-point check*—and the system is safe. One of the most important key insights of backward reachability is the possibility to show the decidability of checking safety properties for some classes of infinite state systems, such as broadcast protocols [33, 27], lossy channel systems [5], timed networks [6], and parametric and distributed systems with global conditions [3, 4]. The main ingredient of the technique for proving decidability of safety is the existence of a well-quasi-ordering over the infinite set of states entailing the termination of backward reachability [1].

**1.1. Array-based systems and symbolic backward reachability.** An *array-based system* (first introduced in [37]) is a generalization of all the classes of infinite state systems mentioned above. Even more, it supports also the specification and verification of algorithms manipulating arrays and fault tolerant systems that are well beyond the paradigms underlying the verification method mentioned above. Roughly, an array-based system is a transition system which updates one (or more) array variable  $\mathbf{a}$ . Being parametric in the structures associated to the indexes and the elements in  $\mathbf{a}$ , the notion of array-based system is quite flexible and allows one the declarative specification of several classes of infinite state systems. For example, consider parametrised systems and the task of specifying their topology: by using no structure at all, indexes are simply identifiers of processes that can only be compared for equality; by using a linear order, indexes are identifiers of processes so that it is possible to distinguish between those on the left or on the right of a process with a particular identifier; by using richer and richer structures (such as trees and graphs), it is possible to specify more and more complex topologies. Similar observations hold also for elements, where it is well-known how to use algebraic structures to specify data structures. Formally, the structure on both indexes and elements is declaratively and uniformly specified by *theories*, i.e. pairs formed by a (first-order) language and a class of (first-order) structures.

On top of the notion of array-based system, it is possible to design a fully symbolic and declarative version of backward reachability for the verification of safety properties where sets of backward reachable states are represented by certain classes of first-order formulae over the signature induced by the theories over the indexes and the elements of the array-based system under consideration. To mechanize this approach, the following three requirements are mandatory:

- (i) the class  $\mathcal{F}$  of (possibly quantified) first-order formulae used to represent sets of states is expressive enough to represent interesting classes of systems and safety properties,
- (ii)  $\mathcal{F}$  is closed under pre-image computation, and
- (iii) the checks for safety and fix-point can be reduced to decidable logical problems (e.g., satisfiability) of formulae in  $\mathcal{F}$ .

Once requirements (i)—(iii) are satisfied, this technique can be seen as a symbolic version of the model checking techniques of [1] revisited in the declarative framework of first-order logic augmented with theories (as first discussed in [37]). Using a declarative framework has several *potential* advantages; two of the most important ones are the following. First, the computation of the pre-image (requirement (ii) above) becomes computationally cheap: we only need to build the formula  $\phi$  representing the (iterated) pre-images of the set of unsafe states and then put the burden of using suitable data structures to represent  $\phi$  on the available (efficient) solver for logical problems encoding safety and fix-point checks. This is in sharp contrast to what is usually done in almost all other approaches to symbolic model checking of infinite state systems, where the computation of the pre-image is computationally very expensive as it requires a substantial process of normalization on the data structure representing the (infinite) sets of states so as to simplify safety and fix-point checks. The second advantage is the possibility to use state-of-the-art Satisfiability Modulo Theories (SMT) solvers, a technology that is showing very good success in scaling up various verification techniques, to support both safety and fix-point checks (requirement (iii) above). Unfortunately, the kind of satisfiability problems obtained in the context of the backward search algorithm requires to cope with (universal) quantifiers and this makes the off-the-shelf use of SMT solvers problematic. In fact, even when using classes of formulae with decidable satisfiability problem, currently available SMT solvers are not yet mature enough to efficiently discharge formulae containing (universal) quantifiers, despite the fact that this problem has recently attracted a lot of efforts (see, e.g., [25, 36, 23]). To alleviate this problem, we have designed a general decision procedure for a class of formulae satisfying requirement (i) above, based on quantifier instantiation (see [37] and Theorem 3.3 below); this allows for an easier way to integrate currently available SMT-solvers in the backward reachability procedure. Interestingly, it is possible to describe the symbolic backward reachability procedure by means of a Tableaux-like calculus which offers a good starting point for implementation. In fact, the main loop of MCMT [41],<sup>1</sup> the model checker for array-based systems that we are currently developing, can be easily understood in terms of the rules of the calculus. The current version of the tool uses Yices [31] as the back-end SMT solver. We have chosen Yices among the many available state-of-the-art solvers because it has scored well in many editions of the SMT-COMP competition and because its lightweight API allowed us to easily embed it in MCMT. An interesting line of future work would be to make the tool parametric with respect to the back-end SMT solver so as to permit the user to select the most appropriate for the problem under consideration.

In our declarative framework, it is also possible to identify sufficient conditions on the theories about indexes and elements of the array-based systems so as to ensure the termination of the symbolic backward reachability procedure. This allows us to derive all the decidability results for the safety problems of the classes of systems mentioned above. Interestingly, the well-quasi-ordering used for the proof of termination can be obtained by using standard model theoretic notions (namely, sub-structures and embeddings) and in conjunction with well-known mathematical results for showing that a binary relation is a well-quasi-order (e.g., Dickson’s Lemma or Kruskal’s Theorem). Contrary to the approach proposed in [1]—where some ingenuity is required, in our framework the definition of well-quasi-order is derived from the class of structures formalizing indexes and elements in a uniform way by using the model-theoretic notions of sub-structure and embedding.

---

<sup>1</sup>The latest available release of the tool with all the benchmarks discussed in this paper (and more) can be downloaded at <http://homes.dsi.unimi.it/~ghilardi/mcmt>.

**1.2. Symbolic backward reachability and Invariant synthesis.** One of the key advantages of backward reachability over other verification methods is to be *goal-directed*; the goal being the set of unsafe states from which pre-images are computed. Despite this, it can unnecessarily explore (large) portions of the symbolic state space of a system which are not required to verify the safety property under consideration. Even worse, in some cases the analysis may not detect a fix-point, thereby causing non-termination. In order to avoid visiting irrelevant parts of the symbolic state space during backward reachability, techniques for analyzing pre-images, over-approximating the set of backward reachable states, and guessing invariants have been devised (see, e.g., [30, 45, 49, 14, 16, 51, 34, 13, 46] to name a few). The success of these techniques depend crucially on the heuristics used to guess the invariants or compute over-approximations. Our approach is similar in spirit to [16], but employs techniques which are specific for our different intended application domains.

Along this line of research, we discuss a technique for interleaving pre-image computation and invariant synthesis which tries to eagerly prune irrelevant parts of the search space. Formally, in the context of the declarative framework described above, our main result about invariant synthesis ensures that the technique *will find an invariant—provided one exists—under suitable hypotheses*, which are satisfied for important classes of array-based systems (e.g., mutual exclusion algorithms or cache coherence protocols). The key ingredient in the proof of the result is again the model-theoretic notion of well-quasi-ordering obtained by applying standard model theoretic notions that already played a key role in showing the termination of the backward reachability procedure. In this case, it allows us to finitely characterize the search space of candidate invariants. Although the technique is developed for array-based systems, we believe that the underlying idea can be adapted to other symbolic approaches to model checking (e.g., [2, 3]).

Although the correctness of our invariant synthesis method is theoretically interesting, its implementation seems to be impractical because of the huge (finite) search space that must be traversed in order to find the desired invariant. In order to make our findings more practically relevant, we study how to integrate invariant synthesis with backward reachability so as to prune the search space of the latter efficiently. To this end, we develop techniques that allow us to analyze a set of backward reachable states and then guess candidate invariants. Such candidate invariants are then proved to be “real” invariants by using a resource bounded variant of the backward reachability procedure and afterwards are used during fix-point checking with the hope that they help pruning the search space by augmenting the chances to detect a fix-point. Two observations are important. First, the bound on the resources of the backward reachability procedure is because we want to obtain invariants in a computationally cheap way. Second, we have complete freedom in the design of the invariant generation techniques as all the candidate invariants are checked to be real invariants before being used by the main backward reachability procedure. As a consequence, (even coarse) abstraction techniques can be used to compute candidate invariants without putting at risk the accuracy of the (un-)safety result returned by the main verification procedure. For concreteness, we discuss two techniques for invariant guessing: both compute over-approximation of the set of backward reachable states. The former, called *index abstraction* (which resembles the technique of [46]), projects away the indexes in the formula used to describe a set of backward reachable states while the latter, called *signature abstraction* (which can be seen as a form of predicate abstraction [44]), projects

away the elements of a sub-set of the array variables by quantifier elimination (if possible). The effectiveness of the proposed invariant synthesis techniques and their integration in the backward reachability procedure must be judged experimentally. Hence, we have implemented them in MCMT and we have performed an experimental analysis on several safety problems translated from available model checkers for parametrised systems (e.g., PFS, Undip, the version of UCLID extended with predicate abstraction) or obtained by the formalizing programs manipulating arrays (e.g., sorting algorithms). The results confirm the viability and the effectiveness of the proposed invariant synthesis techniques either by more quickly finding a fix-point (when the backward reachability procedure alone was already able to find it) or by allowing to find a fix-point (when the backward reachability procedure alone was not terminating).

**How to read the paper.** Given the size of the paper, we identify two tracks for the reader. The former is the ‘symbolic’ track and allows one to focus on the declarative framework, the mechanization of the backward reachability procedure, its combination with invariant synthesis techniques, and its experimental evaluation. The latter is the ‘semantic’ track which goes into the details of the connection between the syntactic characterization of sets of states and the well-quasi-ordering permitting one to prove the termination of the backward reachability procedure and the completeness of invariant synthesis. To some extent, the two tracks can be read independently.

- *Symbolic track.* In Sections 2.1 and 2.3, some preliminary notions underlying the concept of array-based system (Section 3.1) are given. In Section 3, the symbolic version of backward reachability is described, requirements for its mechanization are considered, namely closure under pre-image computation and decidability of safety and fix-point checks (Section 3.2), and its formalization using a Tableaux-like calculus is presented (Section 3.3). In Section 5.1, the notion of safety invariants is introduced, their synthesis and use to prune the search space of the backward reachability procedure is described, and their implementation is considered in Section 6. Particular care has been put in the experimental evaluation of the proposed techniques for invariant synthesis as illustrated in Section 6.4.
- *Semantic track.* In Section 2.2, some notions related to the model theoretic concept of embedding are briefly summarized. In Section 4, it is explained how a pre-order can be defined on sets of states by using the notion of embedding and how this allows us (in case the pre-order is a well-quasi-order) to prove the termination of the backward reachability procedure designed in Section 3. For the sake of completeness, it is also stated that the safety problem for array-based system is undecidable (Section 4.2) and its proof can be found in the Appendix. In Section 5.2, the completeness of an algorithm for invariant synthesis (obtained as the dual of backward reachability) is proved under suitable hypotheses.

In Section 7, we conclude the paper by positioning our work with respect to the state-of-the-art in verification of the safety of infinite state systems and we sketch some lines of future work. For ease of reference, at the end of the paper, we include the table of contents and a figure depicting the two tracks for reading mentioned above.

## 2. FORMAL PRELIMINARIES

We assume the usual syntactic (e.g., signature, variable, term, atom, literal, and formula) and semantic (e.g., structure, truth, satisfiability, and validity) notions of first-order logic (see, e.g., [32]). The equality symbol  $=$  is included in all signatures considered below. A signature is *relational* if it does not contain function symbols and it is *quasi-relational* if its function symbols are all constants. An *expression* is a term, an atom, a literal, or a formula. Let  $\underline{x}$  be a finite tuple of variables and  $\Sigma$  a signature; a  $\Sigma(\underline{x})$ -expression is an expression built out of the symbols in  $\Sigma$  where at most the variables in  $\underline{x}$  may occur free (we will write  $E(\underline{x})$  to emphasize that  $E$  is a  $\Sigma(\underline{x})$ -expression). Let  $\underline{e}$  be a finite sequence of expressions and  $\sigma$  a substitution;  $\underline{e}\sigma$  is the result of applying the substitution  $\sigma$  to each element of the sequence  $\underline{e}$ .

According to the current practice in the SMT literature [52], a *theory*  $T$  is a pair  $(\Sigma, \mathcal{C})$ , where  $\Sigma$  is a signature and  $\mathcal{C}$  is a class of  $\Sigma$ -structures; the structures in  $\mathcal{C}$  are the *models* of  $T$ . Below, we let  $T = (\Sigma, \mathcal{C})$ . A  $\Sigma$ -formula  $\phi$  is *T-satisfiable* if there exists a  $\Sigma$ -structure  $\mathcal{M}$  in  $\mathcal{C}$  such that  $\phi$  is true in  $\mathcal{M}$  under a suitable assignment to the free variables of  $\phi$  (in symbols,  $\mathcal{M} \models \phi$ ); it is *T-valid* (in symbols,  $T \models \varphi$ ) if its negation is *T-unsatisfiable*. Two formulae  $\varphi_1$  and  $\varphi_2$  are *T-equivalent* if  $\varphi_1 \leftrightarrow \varphi_2$  is *T-valid*. The *quantifier-free satisfiability modulo the theory T (SMT(T)) problem* amounts to establishing the *T-satisfiability* of quantifier-free  $\Sigma$ -formulae.

$T$  admits *quantifier elimination* iff given an arbitrary formula  $\varphi(\underline{x})$ , it is always possible to compute a quantifier-free formula  $\varphi'(\underline{x})$  such that  $T \models \forall \underline{x}(\varphi(\underline{x}) \leftrightarrow \varphi'(\underline{x}))$ . Linear Arithmetics, Real Arithmetics, acyclic lists, and enumerated data-type theories (see below) are examples of theories that admit elimination of quantifiers.

A theory  $T = (\Sigma, \mathcal{C})$  is said to be *locally finite* iff  $\Sigma$  is finite and, for every finite set of variables  $\underline{x}$ , there are finitely many  $\Sigma(\underline{x})$ -terms  $t_1, \dots, t_{k_{\underline{x}}}$  such that for every further  $\Sigma(\underline{x})$ -term  $u$ , we have that  $T \models u = t_i$  (for some  $i \in \{1, \dots, k_{\underline{x}}\}$ ). The terms  $t_1, \dots, t_{k_{\underline{x}}}$  are called  *$\Sigma(\underline{x})$ -representative terms*; if they are effectively computable from  $\underline{x}$  (and  $t_i$  is computable from  $u$ ), then  $T$  is said to be *effectively locally finite* (in the following, when we say ‘locally finite’, we in fact always mean ‘effectively locally finite’). If  $\Sigma$  is relational or quasi-relational, then any  $\Sigma$ -theory  $T$  is locally finite.

An important class of theories, ubiquitously used in verification, formalizes enumerated data-types. An *enumerated data-type theory*  $T$  is a theory in a quasi-relational signature whose class of models contains only a single finite  $\Sigma$ -structure  $\mathcal{M} = (M, \mathcal{I})$  such that for every  $m \in M$  there exists a constant  $c \in \Sigma$  such that  $c^{\mathcal{I}} = m$ . For example, enumerated data-type theories can be used to model control locations of processes in parametrised systems (see Example 3.1 below).

**2.1. Case Defined Functions.** In the SMT-LIB format [53], it is possible to use if-then-else constructors when building terms. This may seem to be beyond the realm of first order logic, but in fact these constructors can be easily eliminated in SMT problems. Since case-defined functions (introduced via nested if-then-else constructors) are quite useful for us too, we briefly explain the underlying formal aspects here. Given a theory  $T$ , a *T-partition* is a finite set  $C_1(\underline{x}), \dots, C_n(\underline{x})$  of quantifier-free formulae such that  $T \models \forall \underline{x} \bigvee_{i=1}^n C_i(\underline{x})$  and  $T \models \bigwedge_{i \neq j} \forall \underline{x} \neg(C_i(\underline{x}) \wedge C_j(\underline{x}))$ . A *case-definable extension*  $T' = (\Sigma', \mathcal{C}')$  of a theory  $T = (\Sigma, \mathcal{C})$  is obtained from  $T$  by applying (finitely many times) the following procedure: (i) take a *T-partition*  $C_1(\underline{x}), \dots, C_n(\underline{x})$  together with  $\Sigma$ -terms  $t_1(\underline{x}), \dots, t_n(\underline{x})$ ; (ii) let  $\Sigma'$  be

$\Sigma \cup \{F\}$ , where  $F$  is a “fresh” function symbol (i.e.  $F \notin \Sigma$ ) whose arity is equal to the length of  $\underline{x}$ ; (iii) take as  $\mathcal{C}'$  the class of  $\Sigma'$ -structures  $\mathcal{M}$  whose  $\Sigma$ -reduct is a model of  $T$  and such that  $\mathcal{M} \models \bigwedge_{i=1}^n \forall \underline{x} (C_i(\underline{x}) \rightarrow F(\underline{x}) = t_i(\underline{x}))$ . Thus a case-definable extension  $T'$  of a theory  $T$  contains finitely many additional function symbols, called case-defined functions.

**Lemma 2.1.** *Let  $T'$  be a case-definable extension of  $T$ ; for every formula  $\phi'$  in the signature of  $T'$  it is possible to compute a formula  $\phi$  in the signature of  $T$  such that  $\phi$  and  $\phi'$  are  $T'$ -equivalent.*

*Proof.* It is sufficient to show the claim for an atomic  $\phi'$  containing a single occurrence of a case defined function: if this holds, one can get the general statement by the replacement theorem for equivalent formulae (the procedure must be iterated until all case defined additional function symbols are eliminated). Let  $\phi$  be atomic and let it contain a sub-term of the kind  $F\sigma$  in position  $p$ . Then  $\phi$  is  $T'$ -equivalent to  $\bigvee_i (C_i\sigma \wedge \phi'[t_i\sigma]_p)$ . Here the  $C_i$ 's are the partition formulae for the case definition of  $F$  and the  $t_i$ 's are the related ‘value’ terms; the notation  $\phi'[t_i\sigma]_p$  means the formula obtained from  $\phi'$  by putting  $t_i\sigma$  in position  $p$ .  $\square$

Notice that a case-definable extension  $T'$  of  $T$  is a conservative extension of  $T$ , i.e. formulae in the signature of  $T$  are  $T$ -satisfiable iff they are  $T'$ -satisfiable (this is because, as far as the signature of  $T$  is concerned, the two theories have ‘the same models’). Thus, by Lemma 2.1,  $T$  and  $T'$  are basically the same theory and, by abuse of notation, we shall write  $T$  instead of  $T'$ .

**2.2. Embeddings.** We summarize some basic model-theoretic notions that will be used in Sections 4 and 5 below (for more details, the interested reader is pointed to standard textbooks in model theory, such as [22]).

A  $\Sigma$ -embedding (or, simply, an embedding) between two  $\Sigma$ -structures  $\mathcal{M} = (M, \mathcal{I})$  and  $\mathcal{N} = (N, \mathcal{J})$  is any mapping  $\mu : M \rightarrow N$  among the corresponding support sets satisfying the following three conditions: (a)  $\mu$  is an injective function; (b)  $\mu$  is an algebraic homomorphism, that is for every  $n$ -ary function symbol  $f$  and for every  $a_1, \dots, a_n \in M$ , we have  $f^{\mathcal{N}}(\mu(a_1), \dots, \mu(a_n)) = \mu(f^{\mathcal{M}}(a_1, \dots, a_n))$ ; (c)  $\mu$  preserves and reflects predicates, i.e. for every  $n$ -ary predicate symbol  $P$ , we have  $(a_1, \dots, a_n) \in P^{\mathcal{M}}$  iff  $(\mu(a_1), \dots, \mu(a_n)) \in P^{\mathcal{N}}$ .

If  $M \subseteq N$  and the embedding  $\mu : \mathcal{M} \rightarrow \mathcal{N}$  is just the identity inclusion  $M \subseteq N$ , we say that  $\mathcal{M}$  is a *substructure* of  $\mathcal{N}$  or that  $\mathcal{N}$  is an *superstructure* of  $\mathcal{M}$ . Notice that a substructure of  $\mathcal{N}$  is nothing but a subset of the support set of  $\mathcal{N}$  which is closed under the  $\Sigma$ -operations and whose  $\Sigma$ -structure is inherited from  $\mathcal{N}$  by restriction. In fact, given  $\mathcal{N} = (N, \mathcal{J})$  and  $G \subseteq N$ , there exists the smallest substructure of  $\mathcal{N}$  containing  $G$  in its support set. This is called the substructure *generated by  $G$*  and its support set can be characterized as the set of the elements  $b \in N$  such that  $t^{\mathcal{N}}(\underline{a}) = b$  for some  $\Sigma$ -term  $t$  and some finite tuple  $\underline{a}$  from  $G$  (when we write  $t^{\mathcal{N}}(\underline{a}) = b$ , we mean that  $(\mathcal{N}, \mathbf{a}) \models t(\underline{x}) = y$  for an assignment  $\mathbf{a}$  mapping the  $\underline{a}$  to the  $\underline{x}$  and  $b$  to  $y$ ).

Below, we will make frequent use of the easy—but fundamental—fact that the truth of a universal (resp. existential) sentence is preserved through substructures (resp. through superstructures). A *universal* (resp. *existential*) sentence is obtained by prefixing a string of universal (resp. existential) quantifiers to a quantifier-free formula.



**2.3. A many-sorted framework.** From now on, we use many-sorted first-order logic. All notions introduced above can be easily adapted to a many-sorted framework. **In the rest of the paper, we fix** (i) a theory  $T_I = (\Sigma_I, \mathcal{C}_I)$  whose only sort symbol is INDEX; (ii) a theory  $T_E = (\Sigma_E, \mathcal{C}_E)$  for data whose only sort symbol is ELEM (the class  $\mathcal{C}_E$  of models of this theory is usually a singleton). The **theory  $A_I^E = (\Sigma, \mathcal{C})$  of arrays with indexes in  $T_I$  and elements in  $T_E$**  is obtained as the combination of  $T_I$  and  $T_E$  as follows: INDEX, ELEM, and ARRAY are the only sort symbols of  $A_I^E$ , the signature is  $\Sigma := \Sigma_I \cup \Sigma_E \cup \{-[\_]\}$  where  $[\_]$  has type ARRAY, INDEX  $\rightarrow$  ELEM (intuitively,  $a[i]$  denotes the element stored in the array  $a$  at index  $i$ ); a three-sorted structure  $\mathcal{M} = (\text{INDEX}^{\mathcal{M}}, \text{ELEM}^{\mathcal{M}}, \text{ARRAY}^{\mathcal{M}}, \mathcal{I})$  is in  $\mathcal{C}$  iff  $\text{ARRAY}^{\mathcal{M}}$  is the set of (total) functions from  $\text{INDEX}^{\mathcal{M}}$  to  $\text{ELEM}^{\mathcal{M}}$ , the function symbol  $[\_]$  is interpreted as function application, and  $\mathcal{M}_I = (\text{INDEX}^{\mathcal{M}}, \mathcal{I}_{|\Sigma_I})$ ,  $\mathcal{M}_E = (\text{ELEM}^{\mathcal{M}}, \mathcal{I}_{|\Sigma_E})$  are models of  $T_I$  and  $T_E$ , respectively (here  $\mathcal{I}_{|\Sigma_X}$  is the restriction of the interpretation  $\mathcal{I}$  to the symbols in  $\Sigma_X$  for  $X \in \{I, E\}$ ).

**Notational conventions.** For the sake of brevity, we introduce the following notational conventions:  $d, e$  range over variables of sort ELEM,  $a$  over variables of sort ARRAY,  $i, j, k$ , and  $z$  over variables of sort INDEX. An underlined variable name abbreviates a tuple of variables of unspecified (but finite) length and, if  $\underline{i} := i_1, \dots, i_n$ , the notation  $a[\underline{i}]$  abbreviates the tuple of terms  $a[i_1], \dots, a[i_n]$ . Possibly sub/super-scripted expressions of the form  $\phi(\underline{i}, \underline{e}), \psi(\underline{i}, \underline{e})$  denote **quantifier-free  $(\Sigma_I \cup \Sigma_E)$ -formulae** in which at most the variables  $\underline{i} \cup \underline{e}$  occur. Also,  $\phi(\underline{i}, \underline{t}/\underline{e})$  (or simply  $\phi(\underline{i}, \underline{t})$ ) abbreviates the substitution of the  $\Sigma$ -terms  $\underline{t}$  for the variables  $\underline{e}$ . Thus, for instance,  $\phi(\underline{i}, a[\underline{i}])$  denotes the formula obtained by replacing  $\underline{e}$  with  $a[\underline{i}]$  in a quantifier-free formula  $\phi(\underline{i}, \underline{e})$ .

### 3. BACKWARD REACHABILITY

Following [42], we focus on a particular yet large class of array-based systems corresponding to guarded assignments.

**3.1. Array-based Systems.** A (*guarded assignment*) *array-based (transition) system* (for  $(T_I, T_E)$ ) is a triple  $\mathcal{S} = (a, I, \tau)$  where (i)  $a$  is the *state* variable of sort ARRAY;<sup>2</sup> (ii)  $I(a)$  is the *initial*  $\Sigma(a)$ -formula; and (iii)  $\tau(a, a')$  is the *transition*  $(\Sigma \cup \Sigma_D)(a, a')$ -formula, where  $a'$  is a renamed copy of  $a$  and  $\Sigma_D$  is a finite set of case-defined function symbols not in  $\Sigma_I \cup \Sigma_E$ . Below, we also **assume  $I(a)$  to be a  $\forall^I$ -formula**, i.e. a formula of the form  $\forall \underline{i}. \phi(\underline{i}, a[\underline{i}])$ , and  $\tau(a, a')$  **to be in functional form**, i.e. a *disjunction* of formulae of the form

$$\exists \underline{i} (\phi_L(\underline{i}, a[\underline{i}]) \wedge \forall j a'[j] = F_G(\underline{i}, a[\underline{i}], j, a[j])) \quad (3.1)$$

where  $\phi_L$  is the *guard* (also called the *local* component in [37]), and  $F_G$  is a case-defined function (called the *global* component in [37]). To understand why we say that formulae (3.1) are ‘in functional form’, consider  $\lambda$ -abstraction; then, the sub-formula  $\forall j a'[j] = F_G(\underline{i}, a[\underline{i}], j, a[j])$  can be re-written as  $a' = \lambda j. F_G(\underline{i}, a[\underline{i}], j, a[j])$ . In [37], we adopted a more liberal format for transitions; the format of this paper, however, is sufficient to formalize all relevant examples we met so far. Results in this paper extend in a straightforward way to

<sup>2</sup>For the sake of simplicity, we limit ourselves to array-based systems having just one variable  $a$  of sort ARRAY. All the definitions and results can be easily generalized to the case of several variables of sort ARRAY. In the examples, we will consider cases where more than one variable is required and, in addition, the theory  $T_E$  is many-sorted.

the case in which  $T_E$  is assumed to have quantifier elimination and (3.1) is allowed to have existentially quantified variables ranging over data. This extension is crucial to formalize, e.g., non-deterministic updates or timed networks [20].

Given an array-based system  $\mathcal{S} = (a, I, \tau)$  and a formula  $U(a)$ , (an instance of) the *safety problem* is to establish whether there exists a natural number  $n$  such that the formula

$$I(a_0) \wedge \tau(a_0, a_1) \wedge \cdots \wedge \tau(a_{n-1}, a_n) \wedge U(a_n) \quad (3.2)$$

is  $A_I^E$ -satisfiable. If there is no such  $n$ , then  $\mathcal{S}$  is *safe* (w.r.t.  $U$ ); otherwise, it is *unsafe* since the  $A_I^E$ -satisfiability of (3.2) implies the existence of a run (of length  $n$ ) leading the system from a state in  $I$  to a state in  $U$ . From now on, we **assume**  $U(a)$  **to be a  $\exists^I$ -formula**, i.e. a formula of the form  $\exists \underline{i}. \phi(\underline{i}, a[\underline{i}])$ .

We illustrate the above notions by considering the Mesi cache coherence protocol, taken from the extended version of [2].

**Example 3.1.** Let  $T_I$  be the pure theory of equality and  $T_E$  be the enumerated data-types theory with four constants denoted by the numerals from 1 to 4. Each numeral corresponds to a control location of a cache: 1 to **modified**, 2 to **exclusive**, 3 to **shared**, and 4 to **invalid**.

Initially, all caches are **invalid** and the formula characterizing the set of initial states is  $\forall i. a[i] = 4$ . There are four transitions. In the first (resp. second) transition, a cache in state **invalid** (resp. **shared**) goes to the state **exclusive** and invalidates all the other caches. Formally, these can be encoded with formulae as follows:

$$\begin{aligned} \exists i. (a[i] = 4 \wedge a' = \lambda j. (\text{if } (j = i) \text{ then } 2 \text{ else } 4)) \quad \text{and} \\ \exists i. (a[i] = 3 \wedge a' = \lambda j. (\text{if } (j = i) \text{ then } 2 \text{ else } 4)) \quad . \end{aligned}$$

In the third transition, a cache in state **invalid** goes to the state **shared** and so do all other caches:

$$\exists i. (a[i] = 4 \wedge a' = \lambda j. 3).$$

In the fourth and last transition, a cache in state **exclusive** can move to the state **modified** (the other caches maintain their current state):

$$\exists i. (a[i] = 2 \wedge a' = \lambda j. (\text{if } (j = i) \text{ then } 1 \text{ else } a[j])).$$

To be safe, the protocol should not reach a state in which there is a cache in state **modified** and another cache in state **modified** or in state **shared**. Thus, one can take

$$\exists i_1 \exists i_2. (i_1 \neq i_2 \wedge a[i_1] = 1 \wedge (a[i_2] = 1 \vee a[i_2] = 3))$$

as the unsafety formula. □

The reader with some experience in infinite state model checking may wonder how it is possible to encode in our framework transitions with ‘global conditions,’ i.e. guards requiring a universal quantification over indexes. Indeed, the format (3.1) for transitions is clearly too restrictive for this purpose. However, it is possible to overcome this limitation by using the *stopping failures model* introduced in the literature about distributed algorithms (see, e.g., [47]): according to this model, processes may crash at any time and do not play any role in the rest of the execution of the protocol (they “disappear”). In this model, there is no need to check the universal conditions of a transition, rather the transition is taken and any process not satisfying the global condition is assumed to crash. In this way, we obtain an over-approximation of the original system admitting more runs and any safety

<pre> <b>function</b> BReach(<math>U : \exists^I</math>-formula) 1  <math>P \leftarrow U; B \leftarrow \perp</math>; 2  <b>while</b> (<math>P \wedge \neg B</math> is <math>A_I^E</math>-sat.) <b>do</b> 3      <b>if</b> (<math>I \wedge P</math> is <math>A_I^E</math>-sat.)           <b>then return</b> unsafe; 4      <math>B \leftarrow P \vee B</math>; 5      <math>P \leftarrow Pre(\tau, P)</math>; 6  <b>end</b> 7  <b>return</b> (safe, <math>B</math>);           (a)                 </pre>	<pre> <b>function</b> SInv(<math>U : \exists^I</math>-formula) 1  <math>P \leftarrow ChooseCover(U); B \leftarrow \perp</math>; 2  <b>while</b> (<math>P \wedge \neg B</math> is <math>A_I^E</math>-sat.) <b>do</b> 3      <b>if</b> (<math>I \wedge P</math> is <math>A_I^E</math>-sat.)           <b>then return</b> failure; 4      <math>B \leftarrow P \vee B</math>; 5      <math>P \leftarrow ChooseCover(Pre(\tau, P))</math>; 6  <b>end</b> 7  <b>return</b> (success, <math>\neg B</math>);           (b)                 </pre>
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 1: The basic backward reachability (a) and the invariant synthesis (b) algorithms

certification obtained for this over-approximation is also a safety certification for the original model. Indeed, the converse is not always true and spurious error traces may be obtained. Interestingly, the approximated model can be obtained from the original system by simple syntactical transformations of the formulae encoding the transitions requiring the universal conditions. For more details concerning the implementation of the approximated model in MCMT, the reader is referred to [38]. A more exhaustive discussion of the use of a similar approximated model can be found in [2, 3, 48].

**3.2. Backward Reachable States.** A general approach to solve instances of the safety problem is based on computing the set of backward reachable states. For  $n \geq 0$ , the  $n$ -pre-image of a formula  $K(a)$  is  $Pre^0(\tau, K) := K$  and  $Pre^{n+1}(\tau, K) := Pre(\tau, Pre^n(\tau, K))$ , where

$$Pre(\tau, K) := \exists a'. (\tau(a, a') \wedge K(a')). \quad (3.3)$$

Given  $\mathcal{S} = (a, I, \tau)$  and  $U(a)$ , the formula  $Pre^n(\tau, U)$  describes the set of backward reachable states in  $n$  steps (for  $n \geq 0$ ). At the (end of)  $n$ -th iteration of the loop, the *basic backward reachability algorithm*, depicted in Figure 1 (a), stores in the variable  $B$  the formula  $BR^n(\tau, U) := \bigvee_{i=0}^n Pre^i(\tau, U)$  representing the set of states which are backward reachable from the states in  $U$  in at most  $n$  steps (whereas the variable  $P$  stores the formula  $Pre^{n+1}(\tau, U)$ ). While computing  $BR^n(\tau, U)$ , BReach also checks whether the system is unsafe (cf. line 3, which can be read as ‘ $I \wedge Pre^n(\tau, U)$  is  $A_I^E$ -satisfiable’) or a fix-point has been reached (cf. line 2, which can be read as ‘ $\neg(BR^{n+1}(\tau, U) \rightarrow BR^n(\tau, U))$  is  $A_I^E$ -satisfiable’ or, equivalently, that ‘ $(BR^{n+1}(\tau, U) \rightarrow BR^n(\tau, U))$  is not  $A_I^E$ -valid’). When BReach returns the safety of the system (cf. line 7), the variable  $B$  stores the formula describing the set of states which are backward reachable from  $U$  which is also a fix-point.

Indeed, for BReach (Figure 1 (a)) to be a true (possibly non-terminating) procedure, it is mandatory that (i)  $\exists^I$ -formulae are closed under pre-image computation and (ii) both the  $A_I^E$ -satisfiability test for safety (line 3) and that for fix-point (line 2) are effective.

Concerning (i), it is sufficient to use the following result from [42].<sup>3</sup>

<sup>3</sup>The proposition may be read as the characterization of a weakest liberal pre-condition transformer [29] for array-based systems.

**Proposition 3.2.** *Let  $K(a) := \exists \underline{k} \phi(\underline{k}, a[\underline{k}])$  and  $\tau(a, a') := \bigvee_{h=1}^m \exists \underline{i} (\phi_L^h(\underline{i}, a[\underline{i}]) \wedge a' = \lambda j. F_G^h(\underline{i}, a[\underline{i}], j, a[j]))$ . Then,  $Pre(\tau, K)$  is  $A_I^E$ -equivalent to an (effectively computable)  $\exists^I$ -formula.*

*Proof.* Let  $\tau_h$  be one of the  $m$  disjuncts of  $\tau$ . Using the  $\lambda$ -abstraction formulation and a single  $\beta$ -reduction step, it is clear that  $Pre(\tau_h, K)$  is  $A_I^E$ -equivalent to the following  $\exists^I$ -formula

$$\exists \underline{i} \exists \underline{k}. (\phi_L^h(\underline{i}, a[\underline{i}]) \wedge \phi(\underline{k}, F_G^h(\underline{i}, a[\underline{i}], \underline{k}, a[\underline{k}]))) \quad (3.4)$$

where  $\underline{k}$  is the tuple  $k_1, \dots, k_l$  and  $\phi(\underline{k}, F_G^h(\underline{i}, a[\underline{i}], \underline{k}, a[\underline{k}]))$  is the formula obtained from  $\phi(\underline{k}, a'[\underline{k}])$  by replacing  $a'[k_s]$  with  $F_G^h(\underline{i}, a[\underline{i}], k_s, a[k_s])$ , for  $s = 1, \dots, l$ . Now it is sufficient to eliminate the  $F_G^h$  as shown in Lemma 2.1. As a final step, the existential quantifiers can be moved in front of the disjunction arising from the  $m$  disjuncts  $\tau_1, \dots, \tau_m$ .  $\square$

The proof and the algorithm underlying Proposition 3.2 are quite simple. This is in sharp contrast to most approaches to infinite state model checking available in the literature (e.g., [2, 3]) that use special data structures (such as strings with constraints) to represent sets of states. These special data structures can be considered as normal forms when compared to our formulae. In this respect, our framework is more flexible since—although it can use normal forms (when these can be cheaply computed)—it is not obliged to do so. The drawback is that safety and fix-point checks may become computationally much more expensive. In particular, the bottle-neck is the handling of the quantified variables in the prefix of  $\exists^I$ -formulae which may become quite large at each pre-image computation: notice that the prefix  $\exists \underline{k}$  is augmented with  $\exists \underline{i}$  in (3.4) with respect to  $K$ . This and other issues which are relevant for the implementation of our framework are discussed in [42, 40, 41].

Concerning the mechanization of the safety and fix-point checks (point (ii) above), observe that the formulae involved in the satisfiability checks are  $I \wedge BR^n(\tau, K)$  and  $\neg(BR^n(\tau, U) \rightarrow BR^{n-1}(\tau, U))$ . Since we have closure under pre-image computation, both formulae are of the form  $\exists \underline{a} \exists \underline{i} \forall \underline{j} \psi(\underline{i}, \underline{j}, \underline{a}[\underline{i}], \underline{a}[\underline{j}])$ , where  $\psi$  is quantifier free: we call these sentences  $\exists^{A, I} \forall^I$ -sentences [37].

**Theorem 3.3.** *The  $A_I^E$ -satisfiability of  $\exists^{A, I} \forall^I$ -sentences is decidable if (I)  $T_I$  is locally finite and is closed under substructures<sup>4</sup> and (II) the  $SMT(T_I)$  and  $SMT(T_E)$  problems are decidable. Under the same hypotheses, it holds that an  $\exists^{A, I} \forall^I$ -sentence is  $A_I^E$ -satisfiable iff it is satisfiable in a finite index model (a finite index model is a model  $\mathcal{M}$  in which the set  $INDEX^{\mathcal{M}}$  has finite cardinality).*

A generalization of Theorem 3.3 can be found in the extended version of [37] and is reported in Appendix A (with a proof) to make this paper self-contained. The proof of Theorem 3.3 is the starting point to develop a satisfiability procedure for formulae of the form  $\exists \underline{a} \exists \underline{i} \forall \underline{j} \psi(\underline{i}, \underline{j}, \underline{a}[\underline{i}], \underline{a}[\underline{j}])$  consisting of the following steps: (a) the variables  $\underline{a}, \underline{i}$  are Skolemized away; (b) the variables  $\underline{j}$  are instantiated in all possible ways by using the representative  $\underline{i}$ -terms; (c) the resulting combined problem is purified and an arrangement (i.e. an equivalence class) over the shared index variables is guessed; (d) the positive literals from this arrangement are propagated to the  $T_E$ -literals (this is a variant of the Nelson-Oppen schema adopted in ‘theory connections,’ see [11]); (e) finally, the purified constraints

<sup>4</sup> By this we mean that if  $\mathcal{M}$  is a model of  $T_I$  and  $\mathcal{N}$  is a substructure of  $\mathcal{M}$ , then  $\mathcal{N}$  is a model of  $T_I$  as well.

are passed to the theory solvers for  $T_I$  and  $T_E$ , respectively. From the implementation viewpoint, powerful heuristics are needed [40] to keep the potential combinatorial explosion in step (b) under control. Fortunately, the adoption of a certain format for formulae (called, ‘primitive differentiated,’ see below for details) makes steps (c) and (d) redundant (see [40] for more on this point).

Hypothesis (I) from Theorem 3.3 concerns the *topology* of the system (not the data manipulated by the components of the system) and its intuitive meaning can be easily explained when the signature  $\Sigma_I$  is relational: in that case, local finiteness is guaranteed and closure under substructures says that if some elements are deleted from a model of  $T_I$ , we still get a model of  $T_I$  (i.e. the topology does not change under elimination of elements). For example, Hypothesis (I) is true for (finite) sets, linear orders, graphs, forests, while it does not hold for ‘rings,’ because, after deleting one of their elements, they are no more rings. We emphasize that it is *not possible to weaken* Hypothesis (I) on the theory  $T_I$ . Indeed, it is possible to show that any weakening yields undecidable fragments of the theory of arrays over integers [17] (as it is shown in Appendix A). Furthermore, we observe that Hypothesis (I) is not too restrictive because, as said above, it concerns only the topology of the system. So, for example, the topology of virtually any cache coherence protocol (see Example 3.1) can be formalized by finite sets while that of standard mutual exclusion protocols by linear orders.

We summarize our working hypotheses in the following.

**Assumption 3.4.** *We fix an array-based system  $\mathcal{S} = (a, I, \tau)$  such that the initial formula  $I$  is a  $\forall^I$ -formula, and the transition formula  $\tau(a, a')$  is  $\bigvee_{h=1}^m \tau_h(a, a')$ , where  $\tau_h$  is a formula of the form (3.1) for  $h = 1, \dots, m$ . We suppose that  $\exists$ -formulae are used to describe the set of unsafe states. Finally, we assume that hypotheses (I) and (II) of Theorem 3.3 are satisfied.*

**3.3. Tableaux-like Implementation of Backward Reachability.** A naive implementation of the algorithm in Figure 1 (a) does not scale up. The main problem is the size of the formula  $BR^n(\tau, U)$  which contains many redundant or unsatisfiable sub-formulae. We now discuss how Tableaux-like techniques can be used to circumvent these difficulties. We need one more definition: an  $\exists^I$ -formula  $\exists i_1 \dots \exists i_n \phi$  is said to be *primitive* iff  $\phi$  is a conjunction of literals and is said to be *differentiated* iff  $\phi$  contains as a conjunct the negative literal  $i_k \neq i_l$  for all  $1 \leq k < l \leq n$ . By applying various distributive laws together with the rewriting rules

$$\exists j(i = j \wedge \theta) \rightsquigarrow \theta(i/j) \quad \text{and} \quad \theta \rightsquigarrow (\theta \wedge i = j) \vee (\theta \wedge i \neq j) \quad (3.5)$$

it is always possible to transform every  $\exists^I$ -formula into a disjunction of primitive differentiated ones.

We initialize our tableau with the  $\exists^I$ -formula  $U(a)$  representing the set of unsafe states. The key observation is to revisit the computation of the pre-image as the following inference rule (we use square brackets to indicate the applicability condition of the rule):

$$\frac{K \text{ [} K \text{ is primitive differentiated]}}{Pre(\tau_1, K) \mid \dots \mid Pre(\tau_m, K)} \text{ Prelmg}$$

where  $Pre(\tau_h, K)$  computes the  $\exists^I$ -formula which is  $A_I^E$ -equivalent to the pre-image of  $K$  w.r.t.  $\tau_h$  (this is possible according to the proof of Proposition 3.2).

Since the  $\exists^I$ -formulae labeling the consequents of the rule **Prelmg** may not be primitive and differentiated, we need the following **Beta** rule

$$\frac{K}{K_1 \mid \cdots \mid K_n} \text{Beta}$$

where  $K$  is transformed by applying rewriting rules like (3.5) together with standard distributive laws, in order to get  $K_1, \dots, K_n$  which are primitive, differentiated and whose disjunction is  $A_I^E$ -equivalent to  $K$ .

By repeatedly applying the above rules, it is possible to build a tree whose nodes are labelled by  $\exists^I$ -formulae describing the set of backward reachable states. Indeed, it is not difficult to see that the disjunction of the  $\exists^I$ -formulae labelling all the nodes in the (potentially infinite) tree is  $A_I^E$ -equivalent to the (infinite) disjunction of the formulae  $BR^n(\tau, U)$ , where  $\tau := \bigvee_{h=1}^m \tau_h$ . Indeed, there is no need to fully expand our tree. For example, it is useless to apply the rule **Prelmg** to a node  $\nu$  labelled by an  $\exists^I$ -formula which is  $A_I^E$ -unsatisfiable as all the formulae labelling nodes in the sub-tree rooted at  $\nu$  will also be  $A_I^E$ -unsatisfiable. This observation can be formalized by the following rule closing a branch in the tree (we mark the terminal node of a closed branch by  $\times$ ):

$$\frac{K \ [K \text{ is } A_I^E\text{-unsatisfiable}]}{\times} \text{NotAppl}$$

This rule is effective since  $\exists^I$ -formulae are a subset of  $\exists^{A,I}\forall^I$ -sentences and the  $A_I^E$ -satisfiability of these formulae is decidable by Theorem 3.3.

According to procedure **BReach**, there are two more situations in which we can stop expanding a branch in the tree. One terminates the branch because of the safety test (cf. line 3 of Figure 1 (a)):

$$\frac{K \ [I \wedge K \text{ is } A_I^E\text{-satisfiable}]}{\text{UnSafe}} \text{Safety}$$

Interestingly, if we label with  $\tau_h$  the edge connecting a node labeled with  $K$  with that labeled with  $Pre(\tau_h, K)$  when applying rule **Prelmg**, then the transitions  $\tau_{h_1}, \dots, \tau_{h_e}$  labelling the edges in the branch terminated by **UnSafe** (from the leaf node to the root node) give a *error trace*, i.e. a sequence of transitions leading the array-based system from a state satisfying  $I$  to one satisfying  $U$ . Again, rule **UnSafe** is effective since  $I \wedge K$  is equivalent to an  $\exists^{A,I}\forall^I$ -sentence and its  $A_I^E$ -satisfiability is decidable by Theorem 3.3. The other situation in which one can close a branch corresponds to the fix-point test (cf. line 2 of Figure 1 (a))

$$\frac{K \ [K \wedge \bigwedge \{-K' \mid K' \preceq K\} \text{ is } A_I^E\text{-unsatisfiable}]}{\times} \text{FixPoint}$$

where  $K' \preceq K$  means that  $K'$  is a primitive differentiated  $\exists^I$ -formula labeling a node preceding the node labeling  $K$  (nodes can be ordered according to the strategy for expanding the tree). Once more, this rule is effective since  $K \wedge \bigwedge \{-K' \mid K' \preceq K\}$  can be straightforwardly transformed into an  $\exists^{A,I}\forall^I$ -sentence whose  $A_I^E$ -satisfiability is decidable by Theorem 3.3.

As mentioned above, from the implementation point of view, clever heuristics are needed to reduce the instances that have to be generated for the satisfiability test of Theorem 3.3 and to trivialize the recognition of the unsatisfiable premise of the rule **NotAppl**. In addition, the satisfiability checks required by Rule **FixPoint** should be performed *incrementally* by considering formulae in reverse chronological order (i.e. the pre-images generated later are

added first and those generated early are possibly added later). The interested reader is pointed to [40] for a more exhaustive discussion about these issues.

A final remark is in order. One may think that the main difference between our framework to model checking infinite state systems and other approaches lies just in the technology used for constraint solving; our system, MCMT, uses an SMT solver while other tools (such as PFS [2]) use efficient dedicated algorithms. This is only part of the story. In fact, MCMT usually produces many fewer nodes while visiting the tree whose nodes are labelled with the formulae representing sets of backward reachable states, compared to other systems. This is so because our approach is fully declarative and MCMT *symbolically represents also the topology of the system*, not only the data. The other model checkers use constraints only to represent the data manipulated by the system while the topology is encoded by using an *ad hoc* data structure, which usually requires more effort to represent sets of states. To illustrate this fundamental aspect, we consider a simple (but tricky) example.

**Example 3.5.** Let  $T_I$  be the theory of linear orders and  $T_E$  be an enumerated datatype with 15 constants denoted by the numerals from 1 to 15. Consider the following parametrized system having 7 transitions and 15 control locations:

- the first transition allows process  $i$  to move from location 1 to location 2 provided there is a process  $j$  to the right of  $i$  (i.e.  $i < j$  holds) which is on location 9;
- similarly, the second transition allows process  $i$  to move from location 2 to location 3 provided there is a process  $j$  to the right of  $i$  which is on location 10, and so on (the last transition allows process  $i$  to move from location 7 to location 8 provided there is a process  $j$  to the right of  $i$  which is on location 15).

Initially, all processes are in location 1. We consider the following safety problem: is it possible for a process to reach location 8? The answer is obviously no.

MCMT solves the problem by generating 7 nodes in about 0.02 seconds on a standard laptop. On the contrary, PFS takes about 4 minutes on the same computer and generates thousands of constraints. Why is this so? The point is that tools like PFS do not symbolically represent the system topology and need to specify the relative positions of all the involved processes. In contrast, MCMT can handle partial information like “there exist 7 processes to the right of  $i$  whose locations are from 9 to 15, respectively” just because it is based on a deductive engine, i.e. the SMT solver.

Thus, MCMT represents a fully declarative approach to infinite state model checking that, when coupled with appropriate heuristics, should pave the way to the verification of systems with more and more complex topologies that other tools cannot handle.  $\square$

#### 4. TERMINATION: A SEMANTIC ANALYSIS

Termination of our tableaux calculus (and of the algorithm of Figure 1 (a)) is not guaranteed in general as safety problems are undecidable even when the data structures manipulated by the system are simple (Sec. 4.2). However, it is possible to identify sufficient conditions to obtain termination (Sec. 4.3) which are useful in some applications. We begin by introducing an important definition to be used in this and the following section.

**4.1. Configurations.** A *state* of our array-based system  $\mathcal{S} = (a, I, \tau)$  is a pair  $(s, \mathcal{M})$ , where  $\mathcal{M}$  is a model of  $A_I^E$  and  $s \in \text{ARRAY}^{\mathcal{M}}$ . By recalling the last part of the statement of Theorem 3.3, we can focus on a sub-class of the states (often called configurations) restricting  $\mathcal{M}$  to be a finite index model. Formally, an  $A_I^E$ -*configuration* (or, simply, a *configuration*) is a pair  $(s, \mathcal{M})$  such that  $s$  is an array of a finite index model  $\mathcal{M}$  of  $A_I^E$  ( $\mathcal{M}$  is omitted whenever it is clear from the context). We associate a  $\Sigma_I$ -structure  $s_I$  and a  $\Sigma_E$ -structure  $s_E$  with an  $A_I^E$ -configuration  $(s, \mathcal{M})$  as follows: the  $\Sigma_I$ -structure  $s_I$  is simply the finite structure  $\mathcal{M}_I$ , whereas  $s_E$  is the smallest  $\Sigma_E$ -substructure of  $\mathcal{M}_E$  containing the image of  $s$  (in other words, if  $\text{INDEX}^{\mathcal{M}} = \{c_1, \dots, c_k\}$ , then  $s_E$  is the smallest  $\Sigma_E$ -substructure containing  $\{s(c_1), \dots, s(c_k)\}$ ).

**4.2. Undecidability of the safety problem.** In the general case, safety problems are undecidable. The result is not surprising and we report it in the following for the sake of completeness.

**Theorem 4.1.** *The problem: “given an  $\exists^I$ -formula  $U$ , deciding whether the array-based system  $\mathcal{S}$  is safe w.r.t.  $U$ ” is undecidable (even if  $T_E$  is locally finite).*

The proof consists in a rather straightforward reduction from the reachability problem of Minsky machines. See Appendix A for details.

**4.3. Decidability of the safety problem: sufficient conditions.** A specific feature of array-based systems is that a *partial ordering among configurations* can be defined. This is the key ingredient in establishing the termination of the backward reachability procedure (and thus the decidability of the related safety problem) and characterizing the completeness of invariant synthesis strategies (as it will be shown in Section 5 below).

A *pre-order*  $(P, \leq)$  is a set endowed with a reflexive and transitive relation; an *upset*, also called an *upward closed set*, of such a pre-order is a subset  $U \subseteq P$  such that  $(p \in U$  and  $p \leq q$  imply  $q \in U)$ . An upset  $U$  is *finitely generated* iff it is a finite union of cones, where a *cone* is an upset of the form  $\uparrow p = \{q \in P \mid p \leq q\}$  for some  $p \in P$ . Two elements  $p, q \in P$  are *incomparable* (*equivalent*) if neither (both)  $p \leq q$  nor (and)  $q \leq p$ .

We are ready to define a *pre-order over configurations*. Let  $s, s'$  be configurations: we say that  $s' \leq s$  holds iff there are a  $\Sigma_I$ -embedding  $\mu : s'_I \rightarrow s_I$  and a  $\Sigma_E$ -embedding  $\nu : s'_E \rightarrow s_E$  such that the set-theoretical compositions of  $\mu$  with  $s$  and of  $s'$  with  $\nu$  are equal. This is depicted in the following diagram:

$$\begin{array}{ccc} s'_I & \xrightarrow{\mu} & s_I \\ s' \downarrow & & \downarrow s \\ s'_E & \xrightarrow{\nu} & s_E \end{array}$$

In case  $\mu$  and  $\nu$  are both inclusions, we say that  $s'$  is a *sub-configuration* of  $s$ .

Finitely generated upsets of configurations and  $\exists^I$ -formulae can be used interchangeably under suitable assumptions. Let  $K(a)$  be an  $\exists^I$ -formula; we let  $\llbracket K \rrbracket := \{(s, \mathcal{M}) \mid \mathcal{M} \models K(s)\}$ .

**Proposition 4.2.** *For every  $\exists^I$ -formula  $K(a)$ , the set  $\llbracket K \rrbracket$  is upward closed. For every  $\exists^I$ -formulae  $K_1, K_2$ , we have that  $\llbracket K_1 \rrbracket \subseteq \llbracket K_2 \rrbracket$  iff  $A_I^E \models K_1 \rightarrow K_2$ .*



*Proof.* Let us first show that the set  $\llbracket K \rrbracket$  is upward closed. By using disjunctive normal forms and distributing existential quantifiers over disjunctions, we can suppose—without loss of generality—that  $K(a)$  is of the form  $\exists \underline{i} \phi(\underline{i}, a[\underline{i}])$ , where  $\phi$  is a conjunction of  $\Sigma_I \cup \Sigma_E$ -literals (the general case follows from this one because a union of upsets is an upset). If we also separate  $\Sigma_I$ - and  $\Sigma_E$ -literals, we can suppose that  $\phi(\underline{i}, a[\underline{i}])$  is of the kind  $\phi_I(\underline{i}) \wedge \phi_E(a[\underline{i}])$ , where  $\phi_I$  is a conjunction of  $\Sigma_I$ -literals and  $\phi_E$  is a conjunction of  $\Sigma_E$ -literals. Suppose now that  $(s, \mathcal{M})$  and  $(t, \mathcal{N})$  are configurations such that  $s \leq t$  and  $\mathcal{M} \models K(s)$ : we wish to prove that  $\mathcal{N} \models K(t)$ . From  $\mathcal{M} \models K(s)$ , it follows that there are elements  $\underline{i}$  from  $\text{INDEX}^{\mathcal{M}}$  such that  $\mathcal{M} \models \phi_I(\underline{i}) \wedge \phi_E(s[\underline{i}])$ , i.e. such that  $s_I \models \phi_I(\underline{i})$  and  $s_E \models \phi_E(s[\underline{i}])$  (to infer the latter, recall that the operations  $a[\underline{i}]$  are interpreted as functional applications in our models and also that truth of quantifier free formulae is preserved when considering substructures). Now  $s \leq t$  says that there are embeddings  $\mu : s_I \rightarrow t_I$  and  $\nu : s_E \rightarrow t_E$  such that  $\nu \circ s = t \circ \mu$ . Since truth of quantifier free formulae is preserved when considering superstructures, we get  $t_I \models \phi_I(\mu(\underline{i}))$  and  $t_E \models \phi_E(\nu(s[\underline{i}]))$  (that is,  $t_E \models \phi_E(t(\mu(\underline{i})))$ ) and also  $\mathcal{N} \models \phi_I(\mu(\underline{i})) \wedge \phi_E(t[\mu(\underline{i})])$ , which implies  $\mathcal{N} \models K(t)$ , as desired.

Let us now prove the second claim of the Proposition. That  $A_I^E \models K_1 \rightarrow K_2$  implies  $\llbracket K_1 \rrbracket \subseteq \llbracket K_2 \rrbracket$  is trivial. Suppose conversely that  $A_I^E \not\models K_1 \rightarrow K_2$ , which means that  $K_1(a) \wedge \neg K_2(a)$  is  $A_I^E$ -satisfiable: since this implies that  $K_1(a) \wedge \neg K_2(a)$  is satisfiable in a finite index model of  $A_I^E$  (see Theorem 3.3), we immediately get that  $\llbracket K_1 \rrbracket \not\subseteq \llbracket K_2 \rrbracket$ .  $\square$

Before continuing, we recall the standard model-theoretic notion of Robinson diagrams and some related results (see, e.g., [22] for more details). Let  $\mathcal{M} = (M, \mathcal{I})$  be a  $\Sigma$ -structure which is generated by  $G \subseteq M$ . Let us take a free variable  $x_g$  for every  $g \in G$  and call  $G_x$  the set  $\{x_g \mid g \in G\}$ .<sup>5</sup> The  $\Sigma_G$ -*diagram*  $\delta_{\mathcal{M}}(G)$  of  $\mathcal{M}$  is the set of all  $\Sigma(G_x)$ -literals  $L$  such  $\mathcal{M}, \mathbf{a} \models L$ , where  $\mathbf{a}$  is the assignment mapping  $x_g$  to  $g$ .

The following celebrated result [22] is simple, but nevertheless very powerful and it will be used in the rest of the paper.

**Lemma 4.3** (Robinson Diagram Lemma). *Let  $\mathcal{M} = (M, \mathcal{I})$  be a  $\Sigma$ -structure which is generated by  $G \subseteq M$  and  $\mathcal{N} = (N, \mathcal{J})$  be another  $\Sigma$ -structure. Then, there is a bijective correspondence given by*

$$\mu(g) = \mathbf{a}(x_g) \tag{4.1}$$

(for all  $g \in G$ ) between assignments  $\mathbf{a}$  on  $N$  such that  $\mathcal{N}, \mathbf{a} \models \delta_{\mathcal{M}}(G)$  and  $\Sigma$ -embeddings  $\mu : \mathcal{M} \rightarrow \mathcal{N}$ .

In other words, (4.1) can be used to define  $\mu$  from  $\mathbf{a}$  and conversely. Notice that an embedding  $\mu : \mathcal{M} \rightarrow \mathcal{N}$  is uniquely determined, in case it exists, by the image of the set of generators  $G$ : this is because the fact that  $G$  generates  $\mathcal{M}$  implies (and is equivalent to) the fact that every  $c \in M$  is of the kind  $t^{\mathcal{I}}(g)$ , for some term  $t$  and some  $g$  from  $G$ .

The diagram  $\delta_{\mathcal{M}}(G)$  usually contains infinitely many literals, however there are important cases where we can keep it finite.

**Lemma 4.4.** *Suppose that  $\mathcal{M}$  is a  $\Sigma$ -structure (where  $\Sigma$  is a finite signature), whose support  $M$  is finite; then for every set  $G \subseteq M$  of generators, there are finitely many literals from  $\delta_{\mathcal{M}}(G)$  having all remaining literals of  $\delta_{\mathcal{M}}(G)$  as a logical consequence.*

<sup>5</sup>One may wonder if assuming “countably many variables” is too restrictive since  $G$  may be uncountable. There are two ways to avoid this problem. First, we can use free constants instead of variables (this is the standard solution). Second, we realize that we do not need to consider—in this paper—the case when  $G$  is uncountable since in all our applications,  $G$  is finite.

*Proof.* Choose  $\Sigma(G_x)$ -terms  $t_1, \dots, t_n$  such that (under the assignment  $\mathbf{a} : x_g \mapsto g$ ),  $M$  is equal to the set of the elements assigned by  $\mathbf{a}$  to  $t_1, \dots, t_n$  (this is possible because the elements of  $G$  are generators and  $M$  is finite); we also include the  $x_g$  varying  $g \in G$  among the  $t_1, \dots, t_n$ . We can get the desired finite set  $S$  of literals by taking the set of *atoms* of the form

$$R(t_{i_1}, \dots, t_{i_k}), \quad f(t_{i_1}, \dots, t_{i_k}) = t_{i_{k+1}}$$

(as well as their negations), which are true in  $\mathcal{M}$  under the assignment  $\mathbf{a}$ . In fact, modulo  $S$ , it is easy to see by induction on the structure of the term  $u$  that every  $\Sigma(G_x)$ -term  $u$  is equal to some  $t_i$ ; it follows that every literal from  $\delta_{\mathcal{M}}(G)$  is a logical consequence of  $S$ .  $\square$

Whenever the conditions of the above Lemma are true, we can take a finite conjunction and treat  $\delta_{\mathcal{M}}(G)$  as a single formula: notice that we are allowed to do so whenever  $G$  is finite and  $\mathcal{M}$  is a model of a locally finite theory.

**Proposition 4.5.** *Let  $T_E$  be locally finite. It is possible to effectively associate*

- (i) *an  $\exists^I$ -formula  $K_s$  with every  $A_I^E$ -configuration  $(s, \mathcal{M})$  such that  $\llbracket K_s \rrbracket = \uparrow s$ ;*
- (ii) *a finite set  $\{s_1, \dots, s_n\}$  of  $A_I^E$ -configurations with every  $\exists^I$ -formula  $K$  such that  $K$  is  $A_I^E$ -equivalent to  $K_{s_1} \vee \dots \vee K_{s_n}$ .*

*As a consequence of (i) and (ii), finitely generated upsets of  $A_I^E$ -configurations coincide with sets of  $A_I^E$ -configurations of the kind  $\llbracket K \rrbracket$ , for some  $\exists^I$ -formula  $K$ .*

*Proof.* Ad (i): we take  $G, G'$  to be the support of  $s_I$  and the image of the support of  $s_I$  under the function  $s$ , respectively; clearly  $G$  is a set of generators for  $s_I$  and  $G'$  is a set of generators for  $s_E$ . Let us call the set of variables  $G_x, G'_x$  as  $\underline{i} := \{i_1, \dots, i_n\}$  and  $\underline{e} := \{e_1, \dots, e_n\}$ , respectively. We take  $K_s$  to be

$$\exists \underline{i} (\delta_{s_I}(\underline{i}) \wedge \delta_{s_E}(a_0[\underline{i}])) \tag{4.2}$$

where  $a_0$  is a fresh array variable (in other words, we take the diagrams  $\delta_{s_I}(G), \delta_{s_E}(G')$ , make in the latter the replacement  $\underline{e} \mapsto a_0[\underline{i}]$ , take conjunction, and quantify existentially over the  $\underline{i}$ ). For every configuration  $(t, \mathcal{N})$ , we have that  $t \in \llbracket K_s \rrbracket$  iff  $\delta_{s_I}(\underline{i}) \wedge \delta_{s_E}(a_0[\underline{i}])$  is true in  $\mathcal{N}$  under some assignment  $\mathbf{a}$  mapping the array variable  $a_0$  to  $t$ , that is iff there are embeddings  $\mu : s_I \rightarrow t_I$  and  $\nu : s_E \rightarrow t_E$  as prescribed by Lemma 4.3 (i.e. Robinson Diagram Lemma). These embeddings map the generators  $G$  onto the indexes assigned to the  $\underline{i}$  by  $\mathbf{a}$  and the generators  $G'$  to the elements assigned by  $\mathbf{a}$  to the terms  $a_0[\underline{i}]$ , which means precisely that  $t \circ \mu = \nu \circ s$ . Thus  $t \in \llbracket K_s \rrbracket$  is equivalent to  $s \leq t$ , as desired.

Ad (ii): modulo taking disjunctive normal forms, we can suppose that  $K(a_0)$  is equal to  $\exists \underline{i} \bigvee_k (\phi_k(\underline{i}) \wedge \psi_k(a_0[\underline{i}]))$ , where the  $\phi_k$ 's are  $\Sigma_I$ -formulae, the  $\psi_k$ 's are  $\Sigma_E$ -formulae, and  $\underline{i} := i_1, \dots, i_m$ . Since  $T_I$  is locally finite, we can assume that for every representative  $\underline{i}$ -term  $t$  there is an  $i_s \in \underline{i}$  such that  $t = i_s$  is an  $A_I^E$ -logical consequence of  $\phi_k$ , for all  $k$ : this is achieved by conjoining (just once) equations like  $i_s = t$  with  $\phi_k$  - here the  $i_s$  are new existentially quantified variables and  $t$  is a representative  $\Sigma_I$ -term in which only the original existentially quantified variables occur. In this way, all elements in a substructure generated by  $\underline{i}$  are named explicitly and so are their  $a_0$ -images  $a_0[\underline{i}]$  (otherwise said, modulo  $\phi_k(\underline{i})$ , for every  $\Sigma_I(\underline{i})$ -term  $t$ , we have that  $a_0[t]$  is equal to some of the  $a_0[\underline{i}]$ ).

Now, in a locally finite theory, every quantifier free formula  $\theta$  having at most  $m$  free variables, is equivalent to a disjunction of diagram formulae  $\delta_{\mathcal{M}}(G)$ , where  $\mathcal{M}$  is a substructure of a model of the theory and  $G$  is a set of generators for  $\mathcal{M}$  of cardinality at most  $m$ .<sup>6</sup> If we apply this to both  $T_I$  and  $T_E$ , we get that our  $K(a_0)$  can be rewritten as

$$\bigvee_{\mathcal{A}, \mathcal{B}} \exists \underline{i} (\delta_{\mathcal{A}}(\underline{i}) \wedge \delta_{\mathcal{B}}(a_0[\underline{i}]))$$

where  $\mathcal{A}$  ranges over the  $m$ -generated models of  $T_I$  and  $\mathcal{B}$  over the  $m$ -generated sub-models of  $T_E$  (recall that  $T_I$  is closed under substructures). Every such pair  $(\mathcal{A}, \mathcal{B})$  is either  $A_I^E$ -inconsistent (in case some equality among the generators of  $\mathcal{A}$  is not satisfied by the corresponding generators of  $\mathcal{B}$ ) or it gives rise to a configuration  $a$  such that  $\exists \underline{i} (\delta_{\mathcal{A}}(\underline{i}) \wedge \delta_{\mathcal{B}}(a_0[\underline{i}]))$  is precisely  $K_a$ .  $\square$

The formula  $K_s$  from Proposition 4.5(i) will be called *the diagram formula* for the configuration  $s$ .

The set  $\mathcal{B}(\tau, K)$  of configurations which are backward reachable from the configurations satisfying a given  $\exists^I$ -formula  $K$  is thus an upset, being the union of infinitely many upsets; however, even when the latter are finitely generated,  $\mathcal{B}(\tau, K)$  needs not be so. Under the hypothesis of local finiteness of  $T_E$ , this is precisely what characterizes the termination of the backward reachability procedure.

**Theorem 4.6** ([37]). *Assume that  $T_E$  is locally finite; let  $K$  be an  $\exists^I$ -formula. If  $K$  is safe, then  $\text{BReach}$  in Figure 1 terminates iff  $\mathcal{B}(\tau, K)$  is a finitely generated upset.*<sup>7</sup>

*Proof.* Suppose that  $\mathcal{B}(\tau, K)$  is a finitely generated upset. Notice that

$$\mathcal{B}(\tau, K) = \bigcup_n \llbracket BR^n(\tau, K) \rrbracket,$$

consequently (since we have  $\llbracket BR^0(\tau, K) \rrbracket \subseteq \llbracket BR^1(\tau, K) \rrbracket \subseteq \llbracket BR^2(\tau, K) \rrbracket \subseteq \dots$ ) we have  $\mathcal{B}(\tau, K) = \llbracket BR^n(\tau, K) \rrbracket = \llbracket BR^{n+1}(\tau, K) \rrbracket$  for some  $n$ , which means by the second claim of Proposition 4.2 that  $A_I^E \models BR^n(\tau, K) \leftrightarrow BR^{n+1}(\tau, K)$ , i.e. that the Algorithm halts. Vice versa, if the Algorithm halts, we have  $A_I^E \models BR^n(\tau, K) \leftrightarrow BR^{n+1}(\tau, K)$ , hence  $\llbracket BR^n(\tau, K) \rrbracket = \llbracket BR^{n+1}(\tau, K) \rrbracket = \mathcal{B}(\tau, K)$  and the upset  $\mathcal{B}(\tau, K)$  is finitely generated by Proposition 4.5.  $\square$

To derive a sufficient condition for termination from the Theorem above, we use the notion of a wqo as in [1]. A pre-order  $(P, \leq)$  is a *well-quasi-ordering* (wqo) iff for every sequence

$$p_0, p_1, \dots, p_i, \dots \tag{4.3}$$

of elements from  $P$ , there are  $i < j$  with  $p_i \leq p_j$ .

**Corollary 4.7.**  *$\text{BReach}$  always terminates whenever the pre-order on  $A_I^E$ -configurations is a wqo.*

<sup>6</sup>Since the theory is locally finite, there are finitely many atoms whose free variables are included in a given set of cardinality  $m$ . Maximal conjunctions of literals built on these atoms are either inconsistent (modulo the theory) or satisfiable in an  $m$ -generated substructure of a model of the theory. Because of maximality, these (maximal) conjunctions must be diagrams.

<sup>7</sup>If  $K$  is unsafe, we already know that  $\text{BReach}$  terminates because it detects unsafety.

*Proof.* It is sufficient to show that in a wqo all upsets are finitely generated. This is a well-known fact that can be proved for instance as follows. Let  $U$  be an upset. If  $U$  is empty, then it is finitely generated. Otherwise pick  $p_0 \in U$ , if  $\uparrow p_0 = U$ , clearly  $U$  is finitely generated; otherwise, let  $p_1 \in U \setminus \uparrow p_0$ . At the  $(i+1)$ -th step, either  $U = \uparrow p_0 \cup \dots \cup \uparrow p_i$  and  $U$  is finitely generated, or we can pick  $p_{i+1} \in U$  with  $p_{i+1} \notin \uparrow p_0 \cup \dots \cup \uparrow p_i$ . Since the last alternative sooner or later becomes impossible (because in an infinite sequence like (4.3), we must have  $p_j \in \bigcup_{i < j} \uparrow p_i$  for some  $j$ ), we conclude that  $U$  is finitely generated.  $\square$

Termination of backward reachability for some classes of systems (already considered in the literature) can be obtained from Corollary 4.7; some of these are briefly considered in the example below. Although decidable, many of these cases have very bad computational behavior as only a non-primitive recursive lower bound is known to exist. For the detailed formalization of the classes of systems mentioned in the example below, the interested reader is pointed to the extended version of [37].

**Example 4.8.** We consider three classes of systems for which decidability of the safety problem can be shown by using Corollary 4.7 and well-known results (such as Dickson's Lemma, Higman's Lemma, or Kruskal's theorem; see, e.g., [35] for a survey) for proving that the ordering on configurations is a wqo.

- Take  $T_E$  to be an enumerated data-type theory and  $T_I$  to be the pure theory of equality over the signature  $\Sigma_I = \{=\}$ : the pre-order on  $A_I^E$ -configurations is a wqo by Dickson's Lemma. In fact, if  $T_E$  is the theory of a finite structure with support  $\{e_1, \dots, e_k\}$ , a configuration is uniquely determined by a  $k$ -tuple of integers (counting the number of the  $i$  for which  $a[i] = e_j$  holds) and the configuration ordering is obtained by component-wise comparison. In this setting, one can formalize both cache-coherence [26] (see also Example 3.1) and broadcast protocols [33, 27].
- Take  $T_E$  to be an enumerated data-type theory and  $T_I$  to be the theory of total order: the pre-order on  $A_I^E$ -configurations is a wqo by Higman's Lemma. In fact, if  $T_E$  is the theory of a finite structure with support  $\{e_1, \dots, e_k\}$ , a configuration is uniquely determined by a word on  $\{e_1, \dots, e_k\}$  and the configuration ordering is simply the sub-word relation. In this setting, one can formalize Lossy Channel Systems [5, 50].
- Take  $T_E$  to be the theory of rationals (with the standard ordering relation  $<$ ) and  $T_I$  to be the pure theory of equality over the signature  $\Sigma_I = \{=\}$ : the pre-order on  $A_I^E$ -configurations is a wqo by Kruskal's theorem. In fact, we can represent a configuration  $(s, \mathcal{M})$  as a list  $n_1, \dots, n_k$  of natural numbers (of length  $k$ ): such a list encodes the information that  $s_E$  is a  $k$ -element chain and that  $n_1$  elements from  $s_I$  are mapped by  $s$  into the first element of the chain,  $n_2$  elements from  $s_I$  are mapped by  $s$  into the second element of the chain, etc. If  $w$  is the list for  $s$  and  $v$  is the list for  $s'$ , we have  $s' \leq s$  iff  $w$  is less than or equal component-wise to a sub-word of  $v$ . Termination by Kruskal's theorem is obtained by representing numbers as numerals and by using a binary function symbol  $f$  to encode the precedence (thus, for instance, the list 1,2,2 is represented as  $f(\text{succ}(0), f(\text{succ}(\text{succ}(0)), \text{succ}(\text{succ}(0))))$ ); it is easily seen that, on these terms, the homeomorphic embedding [10] behaves like our configuration ordering.

A final remark is in order. In the model checking literature of infinite state systems, an important property is that of 'monotonicity' [1] (in an appropriate setting, this property is shown to be equivalent to the fact that the pre-image of an upset is still an upset). Such a property is not used in the proofs above as we work symbolically with definable upsets. However, it is possible to formulate it in our framework as follows:

- if  $(s, \mathcal{M}), (s', \mathcal{M}')$ , and  $(t, \mathcal{M})$  are configurations such that  $s \leq s'$  and  $\mathcal{M} \models \tau(s, t)$ , then there exists  $(t', \mathcal{M}')$  such that  $t \leq t'$  and  $\mathcal{M}' \models \tau(s', t')$ .

The proof that such a property holds for transitions in the format (3.1) is easy and left as an exercise to the reader (it basically depends on the fact that truth of existential formulae is preserved by superstructures).

## 5. INVARIANTS SEARCH

It is well-known that invariants are useful for pruning the search space of backward reachability procedures and may help either to obtain or to speed up termination.

**5.1. Safety Invariants.** First of all, we recall the basic notion of safety invariant.

**Definition 5.1.** The  $\forall^I$ -formula  $J(a)$  is a *safety invariant* for the safety problem consisting of the array-based system  $\mathcal{S} = (a, I, \tau)$  and unsafe  $\exists^I$ -formula  $U(a)$  iff the following conditions hold:

- (i)  $A_I^E \models \forall a(I(a) \rightarrow J(a))$ ,
- (ii)  $A_I^E \models \forall a \forall a'(J(a) \wedge \tau(a, a') \rightarrow J(a'))$ , and
- (iii)  $\exists a.(U(a) \wedge J(a))$  is  $A_I^E$ -unsatisfiable.

If we are not given the  $\exists^I$ -formula  $U(a)$  and only conditions (i)–(ii) hold, then  $J(a)$  is said to be an *invariant for  $\mathcal{S}$* .

Checking whether conditions (i), (ii), and (iii) above hold can be reduced, by trivial logical manipulations, to the  $A_I^E$ -satisfiability of  $\exists^{A,I}\forall^I$ -formulae, which is decidable by Theorem 3.3. So, establishing whether a given  $\forall^I$ -formula  $J(a)$  is a safety invariant can be completely automated.

**Property 5.2.** *Let  $U$  be an  $\exists^I$ -formula. If there exists a safety invariant for  $U$ , then the array-based system  $\mathcal{S} = (a, I, \tau)$  is safe with respect to  $U$ .*

*Proof.* For reductio, suppose that there is a safety invariant for  $U$  and the array-based system  $\mathcal{S} = (a, I, \tau)$  is not safe w.r.t.  $U$ . This implies that the formula

$$I(a_0) \wedge \tau(a_0, a_1) \wedge \cdots \wedge \tau(a_{n-1}, a_n) \wedge U(a_n) \quad (5.1)$$

is  $A_I^E$ -satisfiable. By using (i) and (ii) in Definition 5.1, we derive that  $J(a_n) \wedge U(a_n)$  is  $A_I^E$ -satisfiable, in contrast to (iii) in Definition 5.1.  $\square$

Thus, if we are given a suitable safety invariant, Property 5.2 can be used as the basis of the safety invariant method, which turns out to be more powerful than the basic backward reachability procedure in Figure 1 (a).

**Property 5.3.** *Let the procedure BReach in Figure 1(a) terminate on the safety problem consisting of the array-based system  $\mathcal{S} = (a, I, \tau)$  and unsafe formula  $U(a)$ . If BReach returns (safe,  $B$ ), then  $\neg B$  is a safety invariant for  $U$ .*

*Proof.* Suppose that BReach exits the main loop at the  $k$ -th iteration by returning  $B$ ; then  $B$  is  $\bigvee_{i=0}^k \text{Pre}^i(\tau, U)$ ,<sup>8</sup> the formula  $\text{Pre}^{k+1}(\tau, U) \wedge \neg B$  is  $A_I^E$ -unsatisfiable and the formulae  $I \wedge \text{Pre}^i(\tau, U)$  (for  $i = 0, \dots, k$ ) are also  $A_I^E$ -unsatisfiable. The latter means that  $A_I^E \models$

<sup>8</sup>Notice that the disjunction of  $\exists^I$ -formulae is (up to logical equivalence) an  $\exists^I$ -formula, so  $B$  is itself an  $\exists^I$ -formula.

$\forall a(I(a) \rightarrow \neg B(a))$ ; for  $i = 0$  (since  $Pre^0(\tau, U)$  is  $U$ ), we also get that  $\exists a.(U(a) \wedge \neg B(a))$  is  $A_I^E$ -unsatisfiable. To claim that  $\neg B(a)$  is an invariant, we only need to check that  $A_I^E \models \forall a \forall a' (\neg B(a) \wedge \tau(a, a') \rightarrow \neg B(a'))$ , i.e. that  $A_I^E \models \forall a (Pre(\tau, B(a)) \rightarrow B(a))$ , which trivially holds since  $Pre(\tau, B)$  is  $\bigvee_{i=1}^{k+1} Pre^i(\tau, U)$  and hence implies  $Pre^{k+1}(\tau, U) \vee B$  and consequently also  $B$  (recall that  $Pre^{k+1}(\tau, U) \wedge \neg B$  is  $A_I^E$ -unsatisfiable).  $\square$

The converse of Proposition 5.3 does not hold: there might be a safety invariant even when  $BReach$  diverges, as illustrated by the following example.<sup>9</sup>

**Example 5.4.** We consider an algorithm to insert an element  $b[0]$  into a sorted array  $b[1], \dots, b[n]$  (this can be seen as a sub-procedure of the insertion sort algorithm). To formalize this, let  $\Sigma_I$  contain one binary predicate symbol  $S$  and one constant symbol  $0$  and  $T_I$  be the theory whose class of models consists of the substructures of the structure having the naturals as domain, with  $0$  interpreted in the obvious way, and  $S$  interpreted as the graph of the successor function. For the sake of simplicity, we shall use a two-sorted theory for data and two array variables: let  $T_E$  be the two-sorted theory whose class of models consists of the single two-sorted structure given by the Booleans (with the constants  $\top, \perp$  interpreted as true and false, respectively) and the rationals (with the usual ordering relation  $<$ ); the array variable  $a$  is a collection of Boolean flags and the array variable  $b$  is the sorted numerical array where  $b[0]$  should be inserted. The initial  $\forall^I$ -formula is represented as follows:

$$\forall i (a[i] = \perp \leftrightarrow i \neq 0) \wedge \forall i_1, i_2 (S(i_1, i_2) \rightarrow i_1 = 0 \vee b[i_1] \leq b[i_2]),$$

saying that the elements in the array  $b$ , whose corresponding Boolean flag is set to false (namely, all except the one at position 0), are arranged in increasing order. The procedure can be formalized by using just one transition formula in the format 3.1 whose guard and global component are as follows:

$$\begin{aligned} \phi_L(i_1, i_2, a[i_1], a[i_2]) &:= S(i_1, i_2) \wedge a[i_1] = \top \wedge a[i_2] = \perp \wedge b[i_1] > b[i_2] \\ F_G(i_1, i_2, a[i_1], a[i_2], b[i_1], b[i_2], j) &:= \text{if } (j = i_1) \text{ then } \langle \top, b[i_2] \rangle \\ &\quad \text{else if } (j = i_2) \text{ then } \langle \top, b[i_1] \rangle \\ &\quad \text{else } \langle a[j], b[j] \rangle, \end{aligned}$$

which swaps two elements in the array  $b$  if their order is decreasing and sets the Boolean fields appropriately (notice that  $F_G$  updates a pair of array variables whose first component is the new value of  $a$  and second component is the new value of  $b$ ). The obvious correctness property is that there are no two values in decreasing order in the array  $b$  if the corresponding Boolean flags do not allow the transition to fire:

$$\exists i_1, i_2 (S(i_1, i_2) \wedge \neg(a[i_1] = \top \wedge a[i_2] = \perp) \wedge b[i_1] > b[i_2]). \quad (5.2)$$

Unfortunately,  $BReach$  in Figure 1 (a) diverges when applied to (5.2). Fortunately, a safety invariant for (5.2) exists. This can be obtained as follows: run MCMT on the safety problem given by the disjunction of (5.2) and the formula

$$\exists i, j. (S(i, j) \wedge a[i] = \perp \wedge a[j] = \top) \quad (5.3)$$

saying that two adjacent indexes have their Boolean flags set to  $\perp$  and  $\top$ , respectively. The problem is immediately solved by the tool: by Property 5.3, the formula describing the set of backward reachable states is a safety invariant for the safety problem given by the

<sup>9</sup>More significant examples having a similar behavior can be found in the MCMT distribution.

disjunction of (5.2) and (5.3), hence *a fortiori* also for the safety formula (5.2) alone. In this case, formula (5.3) has been found manually; however, MCMT *can find it without user intervention* as soon as its invariant synthesis capabilities are activated by suitable command line options. The combination of automatic invariant search and backward reachability will be the main subject of Section 6.1 below.  $\square$

It is interesting to rephrase the conditions of Definition 5.1 in terms of configurations as this paves the way to characterize the completeness of our invariant synthesis method as will be shown below.

**Lemma 5.5.** *Let  $J$  be a  $\forall^I$ -formula; the conditions (i), (ii), and (iii) of Definition 5.1 are equivalent to the following three conditions on (sets of) configurations:*

$$\llbracket I \rrbracket \cap \llbracket H \rrbracket = \emptyset \quad (5.4)$$

$$\llbracket \text{Pre}(\tau, H) \rrbracket \subseteq \llbracket H \rrbracket \quad (5.5)$$

$$\llbracket U \rrbracket \subseteq \llbracket H \rrbracket, \quad (5.6)$$

where  $H$  is the  $\exists^I$ -formula which is logically equivalent to the negation of  $J$ .

*Proof.* For (5.4), we have:

$$\begin{aligned} \text{(i) of Def. 5.1} &\Leftrightarrow A_I^E \models \forall a.(I(a) \rightarrow J(a)) \Leftrightarrow \\ &\neg \forall a.(I(a) \rightarrow J(a)) \text{ is } A_I^E\text{-unsat.} \Leftrightarrow \\ &\exists a.(I(a) \wedge \neg J(a)) \text{ is } A_I^E\text{-unsat.} \Leftrightarrow \\ &\exists a.(I(a) \wedge H(a)) \text{ is } A_I^E\text{-unsat.} \Leftrightarrow \llbracket I \rrbracket \cap \llbracket H \rrbracket = \emptyset. \end{aligned}$$

For (5.5), we have:

$$\begin{aligned} \text{(ii) of Def. 5.1} &\Leftrightarrow A_I^E \models \forall a, a'.(J(a) \wedge \tau(a, a') \rightarrow J(a')) \Leftrightarrow \\ &\exists a, a'.\neg(J(a) \wedge \tau(a, a') \rightarrow J(a')) \text{ is } A_I^E\text{-unsat.} \Leftrightarrow \\ &\exists a, a'.(J(a) \wedge \tau(a, a') \wedge \neg J(a')) \text{ is } A_I^E\text{-unsat.} \Leftrightarrow \\ &\exists a.(J(a) \wedge \exists a'.(\tau(a, a') \wedge \neg J(a'))) \text{ is } A_I^E\text{-unsat.} \Leftrightarrow \\ &\exists a.(J(a) \wedge \exists a'.(\tau(a, a') \wedge H(a'))) \text{ is } A_I^E\text{-unsat.} \Leftrightarrow \\ &\exists a.(J(a) \wedge \text{Pre}(\tau, H)(a)) \text{ is } A_I^E\text{-unsat.} \Leftrightarrow \\ &A_I^E \models \forall a.(\neg J(a) \vee \neg \text{Pre}(\tau, H)(a)) \Leftrightarrow \\ &A_I^E \models \forall a.(H(a) \vee \neg \text{Pre}(\tau, H)(a)) \Leftrightarrow \\ &A_I^E \models \forall a.(\text{Pre}(\tau, H)(a) \rightarrow H(a)) \Leftrightarrow \llbracket \text{Pre}(\tau, H) \rrbracket \subseteq \llbracket H \rrbracket. \end{aligned}$$

For (5.6), we have:

$$\begin{aligned} \text{(iii) of Def. 5.1} &\Leftrightarrow \exists a.(U(a) \wedge J(a)) \text{ is } A_I^E\text{-unsat.} \Leftrightarrow \\ &\exists a.\neg(\neg U(a) \vee \neg J(a)) \text{ is } A_I^E\text{-unsat.} \Leftrightarrow \\ &A_I^E \models \forall a.(U(a) \rightarrow \neg J(a)) \Leftrightarrow \\ &A_I^E \models \forall a.(U(a) \rightarrow H(a)) \Leftrightarrow \llbracket U \rrbracket \subseteq \llbracket H \rrbracket. \end{aligned}$$

$\square$

**5.2. Invariant Synthesis.** The main difficulty to exploit Property 5.2 is to find suitable  $\forall^I$ -formulae satisfying conditions (i)—(iii) of Definition 5.1. Unfortunately, the set of  $\forall^I$ -formulae which are candidates to become safety invariants is infinite. Such a search space can be dramatically restricted when  $T_E$  is locally finite, although it is still infinite because there is no bound on the length of the universally quantified prefix. From a technical point of view, we need to develop some preliminary results.

First, we give a closer look to the *equivalence* relation among configurations: we recall that  $s$  is equivalent to  $t$  (written  $s \approx t$ ) iff  $s \leq t$  and  $t \leq s$ .

**Proposition 5.6.** *We have that  $s \approx t$  holds iff there are a  $\Sigma_I$ -isomorphism  $\mu$  and a  $\Sigma_E$ -isomorphism  $\nu$  such that the set-theoretical compositions of  $\mu$  with  $s$  and of  $s'$  with  $\nu$  are equal.<sup>10</sup> The situation is depicted in the following diagram:*

$$\begin{array}{ccc} s'_I & \xrightarrow{\mu} & s_I \\ s' \downarrow & & \downarrow s \\ s'_E & \xrightarrow{\nu} & s_E \end{array}$$

*Proof.* The implication ‘ $\Leftarrow$ ’ is straightforward and thus we detail only ‘ $\Rightarrow$ ’ in the following. The supports of  $s_I$  and of  $t_I$  are finite, hence the existence of embeddings  $s_I \xrightarrow{\mu_1} t_I \xrightarrow{\mu_2} s_I$  means (for cardinality reasons) that  $\mu_1, \mu_2$  are bijections, hence isomorphisms. Since the images of  $s$  and  $t$  are finite sets of generators for  $s_E$  and  $t_E$ , respectively, we have embeddings  $s_E \xrightarrow{\nu_1} t_E \xrightarrow{\nu_2} s_E$  mapping generators into generators: again, for cardinality reasons,  $\nu_1, \nu_2$  restrict to bijections among generators, which means that they are isomorphisms.  $\square$

**Definition 5.7.** A *basis* for a finitely generated upset  $S$  (resp., for an  $\exists^I$ -formula  $K$ ) is a minimal finite set  $\{s_1, \dots, s_n\}$  such that  $S$  (resp.,  $\llbracket K \rrbracket$ ) is equal to  $\uparrow s_1 \cup \dots \cup \uparrow s_n$ .

It is easy to see that two bases for the same upset are essentially the same, in the sense that *they are formed by pairwise equivalent configurations*. Suppose in fact that  $\{s_1, \dots, s_n\}$  and  $\{s'_1, \dots, s'_m\}$  are two bases for the same upset. Then for every  $s_i$  there exists  $s'_j$  such that  $s'_j \leq s_i$ ; however, there is also  $s_k$  with  $s_k \leq s'_j$  (because  $\{s_1, \dots, s_n\}$  is a basis) and by minimality it follows that  $s_i = s_k$ , which means that  $s_i$  and  $s'_j$  are equivalent. Thus each member of a basis is equivalent to a member of the other (and to a unique one by minimality again) and vice versa; in particular, we also have that  $m = n$ .

**Lemma 5.8.** *Suppose  $T_E$  is locally finite. A configuration  $s$  belongs to a basis for an  $\exists^I$ -formula  $K$  iff  $s \in \llbracket K \rrbracket$  and for every  $s'$  ( $s' \leq s$  and  $s' \in \llbracket K \rrbracket$ ) imply that  $s \approx s'$ .*

*Proof.* Let  $B$  be a basis for  $K$  and let also  $s \in B$ ,  $s' \leq s$  and  $s' \in \llbracket K \rrbracket$ ; then  $s'$  is bigger than some configuration from  $B$ , which must be  $s$ , because elements from  $B$  are incomparable:  $s \approx s'$  follows immediately. Conversely, suppose that  $s \in \llbracket K \rrbracket$  and for every  $s'$ ,  $s' \leq s$  and  $s' \in \llbracket K \rrbracket$  imply that  $s \approx s'$ . Since  $T_E$  is locally finite,  $K$  has a basis  $B$  (this can

<sup>10</sup>Notice that, since the image of  $s$  is a set of generators for  $s_E$ , it is not difficult to see that  $\nu$  is uniquely determined from  $\mu$  (i.e., given  $\mu$ , there might be no  $\nu$  such that the square commutes, but in case one such exists, it is unique). Observe also that, if  $s$  comes from the finite index model  $\mathcal{M}$  and  $t$  comes from the finite index model  $\mathcal{N}$ , the fact that  $s \approx t$  holds does not mean that  $\mathcal{M}$  and  $\mathcal{N}$  are isomorphic: their  $\Sigma_I$ -reducts are  $\Sigma_I$ -isomorphic, but their  $\Sigma_E$ -reducts need not be  $\Sigma_E$ -isomorphic (only the  $\Sigma_E$ -substructures  $s_E$  and  $t_E$  are  $\Sigma_E$ -isomorphic).



be immediately deduced from Proposition 4.5(ii)). We have  $b \leq s$  (and also  $s \approx b$ ) for some  $b$  from  $B$ : it is now clear that we can get another basis for  $K$  by replacing in  $B$  the configuration  $b$  with  $s$ .  $\square$

Our goal is to integrate the safety invariant method into the basic Backward Reachability algorithm of Figure 1(a). To this end, we introduce the notion of ‘sub-reachability.’

**Definition 5.9** (Subreachable configurations). Suppose  $T_E$  is locally finite and let  $s$  be a configuration. A *predecessor* of  $s$  is any  $s'$  that belongs to a basis for  $Pre(\tau, K_s)$  (see Proposition 4.5 for the definition of  $K_s$ ). Let  $s, s'$  be configurations:  $s$  is *sub-reachable* from  $s'$  iff there exist configurations  $s_0, \dots, s_n$  such that (i)  $s_0 = s$ , (ii)  $s_n = s'$ , and (iii) either  $s_{i-1} \leq s_i$  or  $s_{i-1}$  is a predecessor of  $s_i$ , for each  $i = 1, \dots, n$ . If  $K$  is an  $\exists^I$ -formula,  $s$  is *sub-reachable from  $K$*  iff  $s$  is sub-reachable from some  $s'$  taken from a basis of  $K$ .

The following is the main technical result of this section.

**Theorem 5.10.** *Let  $T_E$  be locally finite. If there exists a safety invariant for  $U$ , then there are finitely many  $A_I^E$ -configurations  $s_1, \dots, s_k$  which are sub-reachable from  $U$  and such that  $\neg(K_{s_1} \vee \dots \vee K_{s_k})$  is also a safety invariant for  $U$ .*

*Proof.* Our goal is to replace an  $\exists^I$ -formula  $H$  satisfying the three conditions of Lemma 5.5 with an  $\exists^I$ -formula  $L$  whose negation is still a safety invariant for  $U$  and whose basis is formed by configurations which are all sub-reachable from  $U$ . To this end, we consider a function  $\gamma(S)$  where  $S$  is an  $\exists^I$ -formula such that  $\llbracket S \rrbracket \subseteq \llbracket H \rrbracket$ : the function  $\gamma(S)$  returns an  $\exists^I$ -formula  $K_{a_1} \vee \dots \vee K_{a_n}$ , where  $\{a_1, \dots, a_n\} \subseteq \llbracket H \rrbracket$  is a minimal set of configurations taken from a basis of  $H$  such that  $\llbracket S \rrbracket \subseteq \uparrow a_1 \cup \dots \cup \uparrow a_n$ . (Notice that this implies that  $\{a_1, \dots, a_n\}$  is a basis of  $\gamma(S)$  and  $\llbracket S \rrbracket \subseteq \llbracket \gamma(S) \rrbracket$ .)<sup>11</sup>

Now, define the following sequence of  $\exists^I$ -formulae  $L_j$ : (i)  $L_0 := \gamma(U)$  and (ii)  $L_{i+1} := L_i \vee \gamma(Pre(\tau, L_i))$ . (The definition is well given because  $\llbracket L_i \rrbracket \subseteq \llbracket H \rrbracket$  is a consequence of (5.6) and (5.5).) What remains to be shown is that the sequence becomes stable and its fix-point is the desired  $L$ , i.e. a safety invariant for  $U$  whose basis is formed by configurations which are sub-reachable from  $U$ .

We first show, by induction on  $k$ , that every configuration  $b$  that belongs to a basis of  $L_k$  is sub-reachable from  $U$ :

- if  $k = 0$ , we have that  $\{a_1, \dots, a_n\}$  is a minimal set of configurations taken from a basis of  $H$  such that  $\llbracket U \rrbracket \subseteq \uparrow a_1 \cup \dots \cup \uparrow a_n$  and  $b = a_j$  for some  $j = 1, \dots, n$ . By minimality, there is  $s$  from a basis of  $U$  such that  $s \notin \uparrow a_1 \cup \dots \cup \uparrow a_{j-1} \cup \uparrow a_{j+1} \cup \dots \cup \uparrow a_n$ , which means that  $s \in \uparrow a_j$ , that is  $a_j \leq s$  and  $a_j = b$  is sub-reachable from  $U$ .
- Suppose now  $k = i + 1 > 0$ . A basis for  $L_i \vee \gamma(Pre(\tau, L_i))$  is obtained by joining two bases—one for  $L_i$  and one for  $\gamma(Pre(\tau, L_i))$ —and then by discarding non-minimal elements. As a consequence, if  $b$  is in a basis for  $L_k$ , then  $b$  is either in a basis for  $L_i$  or in a basis for  $\gamma(Pre(\tau, L_i))$  (or in both). In the former case, we just apply induction. If  $b$  is in a basis for  $\gamma(Pre(\tau, L_i))$ , the same argument used in the case  $k = 0$  shows that  $b \leq s$  for an  $s$  that belongs to a basis for  $Pre(\tau, L_i)$ . Now, if  $c_{i1}, \dots, c_{ik_i}$  is a basis of  $L_i$ , the formula  $Pre(\tau, L_i)$  is  $A_I^E$ -equivalent to the disjunction of the  $Pre(\tau, K_{c_{ij}})$  and

<sup>11</sup>There might be many functions  $\gamma$  satisfying the above specification, we just take one of them. This can be done (by choice axiom) because, given  $S$  such that  $\llbracket S \rrbracket \subseteq \llbracket H \rrbracket$ , there always exists a minimal set of configurations  $\{a_1, \dots, a_n\}$  taken from a basis of  $H$  such that  $\llbracket S \rrbracket \subseteq \uparrow a_1 \cup \dots \cup \uparrow a_n$  (just take any basis for  $H$  and throw out configurations from it until minimality is acquired).

consequently  $s$  must be in a basis of one of the latter (that is,  $s$  is a predecessor of some  $c_{ij}$ ); since the  $c_{ij}$  are sub-reachable by induction hypothesis and  $b \leq s$ , the definition of sub-reachability guarantees that  $b$  is sub-reachable from  $U$ .

The increasing chain

$$\llbracket L_0 \rrbracket \subseteq \llbracket L_1 \rrbracket \subseteq \dots$$

becomes stationary, because at each step only configurations from a basis of  $H$  can be added and bases are (unique and) finite by definition. Thus, we have  $\llbracket L_i \rrbracket = \llbracket L_{i+1} \rrbracket$  for some  $i$ : let  $L$  be  $L_i$  for such  $i$ .

The fact that  $L$  is a safety invariant is straightforward: condition  $\llbracket I \rrbracket \cap \llbracket L \rrbracket = \emptyset$  follows from (5.4) and the fact that  $\llbracket L \rrbracket \subseteq \llbracket H \rrbracket$ , whereas conditions  $\llbracket U \rrbracket \subseteq \llbracket L \rrbracket$  and  $\llbracket \text{Pre}(\tau, L) \rrbracket \subseteq \llbracket L \rrbracket$  follow directly from the above definitions of  $L_0$  and  $L_{i+1}$  (we have  $\llbracket U \rrbracket \subseteq \llbracket \gamma(U) \rrbracket = \llbracket L_0 \rrbracket \subseteq \llbracket L \rrbracket$  and for all  $i \geq 0$ ,  $\llbracket \text{Pre}(\tau, L_i) \rrbracket \subseteq \llbracket \gamma(\text{Pre}(\tau, L_i)) \rrbracket \subseteq \llbracket L_{i+1} \rrbracket \subseteq \llbracket L \rrbracket$ ).  $\square$

The intuition underlying the theorem is as follows. Let us call ‘finitely representable’ an upset which is of the kind  $\llbracket K \rrbracket$  for some  $\exists^I$ -formula  $K$  and let  $B$  be the set of backward reachable states. Usually  $B$  is infinite and it is finitely representable only in special cases (e.g., when the configuration ordering is a wqo). Nevertheless, it may sometimes exist a set  $B' \supseteq B$  which is finitely representable and whose complement is an invariant of the system. Theorem 5.10 ensures us to find such a  $B'$ , if any exists. This is the case of Example 5.4 where not all configurations satisfying (5.3) are in  $B$  and  $B$  must be enlarged to encompass such configurations too (only in this way it becomes finitely representable, witness the fact that backward reachability diverges).

In practice, Theorem 5.10 suggests the following procedure to find the super-set  $B'$ . At each iteration of **BReach**, the algorithm represents symbolically in the variable  $B$  the configurations which are backward reachable in  $n$  steps; before computing the next pre-image of  $B$ , non deterministically replace some of the configurations in a basis of  $B$  with some sub-configurations and update  $B$  by a symbolic representation of the upset obtained in this way. As a consequence, if an invariant exists, we are guaranteed to find it; otherwise, the process may diverge. Notice that (in the local finiteness hypothesis for  $T_E$ ) the search space of the configurations which are sub-reachable in  $n$  steps is finite, although this search space is infinite if no bound on  $n$  is fixed. To illustrate, (5.3) in Example 5.4 contains some sub-reachable only configurations. This shows that sub-reachability is crucial for Theorem 5.10 to hold.

The algorithm sketched above can be refined further so as to obtain a completely symbolic method working with formulae without resorting to configurations. The key idea to achieve this is to rephrase in a symbolic setting the relevant notions concerning sub-reachability. However, this goal is best achieved incrementally as there are some subtle aspects to take care of. The starting point is the following observation. It is not possible to characterize the fact that a configuration  $(s, \mathcal{M})$  is part of a basis for an  $\exists^I$ -formula  $\exists \underline{i} \phi(\underline{i}, a[\underline{i}])$  by using another  $\exists^I$ -formula (a universal quantifier is needed to express the suitable minimality requirement). Instead, we shall characterize by an  $\exists^I$ -formula the fact that a tuple satisfying  $\phi(\underline{i}, a[\underline{i}])$  generates a submodel which is a configuration belonging to a basis (see Lemma 5.11 below). Notice that the simple fact that the tuple satisfies  $\phi$  is not sufficient alone: for instance, only pairs formed by *identical* elements satisfying  $a[i_1] = a[i_2]$  generate a configuration in a basis (tuples formed by pairs of different elements are not

minimal). To generalize this, we introduce the following abbreviation:

$$Min(\phi, a, \underline{i}) := \phi(\underline{i}, a[\underline{i}]) \wedge \bigwedge_{\sigma} \left( \phi(\underline{i}\sigma, a[\underline{i}\sigma]) \rightarrow \bigwedge_{i \in \underline{i}} \bigvee_t (t\sigma = i) \right) \quad (5.7)$$

where  $\phi(\underline{i}, a[\underline{i}])$  is a quantifier-free formula,  $t$  ranges over representative  $\Sigma_I(\underline{i})$ -terms, and  $\sigma$  ranges over the substitutions with domain  $\underline{i}$  and co-domain included in the set of representative  $\Sigma_I(\underline{i})$ -terms. The following lemma gives a semantic characterization of  $Min(\phi, a, \underline{i})$ .

**Lemma 5.11.** *Consider an  $\exists^I$ -formula  $K \equiv \exists \underline{i} \phi(\underline{i}, a[\underline{i}])$ , an  $A_I^E$ -model  $\mathcal{M}$ , and a variable assignment  $\mathbf{a}$  in  $\mathcal{M}$  such that  $(\mathcal{M}, \mathbf{a}) \models \phi(\underline{i}, a[\underline{i}])$ . We have that  $(\mathcal{M}, \mathbf{a}) \models Min(\phi, a, \underline{i})$  iff the configuration  $s$  obtained by restricting  $\mathbf{a}(a)$  to the  $\Sigma_I$ -substructure generated by the  $\mathbf{a}(\underline{i})$ 's belongs to a basis of  $K$ .<sup>12</sup>*

*Proof.* Suppose that  $(\mathcal{M}, \mathbf{a}) \models Min(\phi, a, \underline{i})$  (for simplicity, we shall directly call  $\underline{i}, a$  the elements assigned by  $\mathbf{a}$  to  $\underline{i}, a$ , respectively). By Proposition 5.6 and Lemma 5.8, it is sufficient to show the following. Consider  $s' \leq s$  such that  $s' \in \llbracket K \rrbracket$ : we show that the embeddings  $\mu, \nu$  witnessing the relation  $s' \leq s$  and making the diagram

$$\begin{array}{ccc} s'_I & \xrightarrow{\mu} & s_I \\ s' \downarrow & & \downarrow s \\ s'_E & \xrightarrow{\nu} & s_E \end{array}$$

to commute are isomorphisms (in fact, it is sufficient to show only that  $\mu$  is bijective, because the images of  $s'$  and  $s$  are  $\Sigma_E$ -generators and the square commutes). Without loss of generality, we can assume that  $\mu$  is an inclusion; the domain of  $s'$  is then formed by elements of the form  $t^{(\mathcal{M}, \mathbf{a})}$  for suitable (representative)  $\Sigma_I(\underline{i})$ -terms  $t$  and the fact that  $s' \in \llbracket K \rrbracket$  means then that  $(\mathcal{M}, \mathbf{a}) \models \phi(\underline{i}\sigma, a[\underline{i}\sigma])$  holds for a substitution  $\sigma$  whose domain is  $\underline{i}$  and whose range is contained into the set of those representative  $\Sigma_I(\underline{i})$ -terms  $u$  such that  $u^{(\mathcal{M}, \mathbf{a})}$  is in the support of  $s'_I$ . Since  $(\mathcal{M}, \mathbf{a}) \models Min(\phi, a, \underline{i})$  holds, for every  $i \in \underline{i}$  there is a representative  $\Sigma_I(\underline{i})$ -term  $t$  such that  $(\mathcal{M}, \mathbf{a}) \models t\sigma = i$  holds. The latter means that  $i$  is in the support of  $s'_I$ , hence the inclusion  $\mu$  is onto.

Conversely, if  $s$  belongs to a basis of  $K$ , then there is no  $s' \leq s$  is in  $\llbracket K \rrbracket$ , unless  $s'$  is equivalent to  $s$ , by Lemma 5.8. Suppose that  $(\mathcal{M}, \mathbf{a}) \models \phi(\underline{i}\sigma, a[\underline{i}\sigma])$  holds for a substitution  $\sigma$  whose domain is  $\underline{i}$  and whose range is included into the set of representative  $\Sigma_I(\underline{i})$ -terms. For reductio, suppose that  $(\mathcal{M}, \mathbf{a}) \models t\sigma = i$  does not hold for some  $i \in \underline{i}$  and all representative  $\Sigma_I(\underline{i})$ -terms  $t$ ; we can restrict the array  $a$  to the  $\Sigma_I$ -substructure given by the elements of the kind  $t\sigma^{(\mathcal{M}, \mathbf{a})}$ , thus getting a configuration  $s' \leq s$  such that  $s' \in \llbracket K \rrbracket$ . Since the finite support of  $s'_I$  has smaller cardinality than the support of  $s_I$  (because  $\mathbf{a}(i)$  does not belong to it), we cannot have  $s' \approx s$ , a contradiction!  $\square$

<sup>12</sup>To make the statement of the lemma precise, one should define not just  $s$  but also the finite index model where  $s$  is taken from. In detail, we take the  $A_I^E$ -model  $\mathcal{N}$  whose  $\Sigma_I$ -reduct is the restriction of  $\mathcal{M}_I$  to the  $\Sigma_I$ -substructure generated by the  $\mathbf{a}(\underline{i})$ 's and whose  $\Sigma_E$ -reduct is equal to  $\mathcal{M}_E$ . In this model, we can define the array  $s$  to be the restriction of  $\mathbf{a}(a)$  to  $\text{INDEX}^{\mathcal{N}} \subseteq \text{INDEX}^{\mathcal{M}}$ . The pair  $(s, \mathcal{N})$  is now a configuration in the sense defined in Section 5.2.

**Remark 5.12.** We identify conditions under which it is trivial to compute  $Min(\phi, a, \underline{i})$ . Besides being an interesting observation *per se*, it will be used later in this section to illustrate simple and useful examples of the key notion of cover (see Example 5.17 below). If (as it often happens in applications) the signature  $\Sigma_I$  is relational and the formula  $\phi(\underline{i}, a[\underline{i}])$  is differentiated,  $Min(\phi, a, \underline{i})$  is  $A_I^E$ -equivalent to  $\phi(\underline{i}, a[\underline{i}])$ : this is because only variable permutations can be consistently taken into consideration as the  $\sigma$ 's in formula (5.7), so that the  $t\sigma$ 's are precisely the  $\underline{i}$ 's.

**Corollary 5.13.** *Consider an  $\exists^I$ -formula  $K := \exists \underline{i} \phi(\underline{i}, a[\underline{i}])$  and a configuration  $(s, \mathcal{M})$ ; if  $s$  belongs to a basis for  $K$ , then  $(\mathcal{M}, \mathbf{a}) \models \phi(\underline{i}, a[\underline{i}]) \rightarrow Min(\phi, a, \underline{i})$  holds for all  $\mathbf{a}$  such that  $\mathbf{a}(a) = s$ .*

*Proof.* If  $(\mathcal{M}, \mathbf{a}) \models \phi(\underline{i}, a[\underline{i}])$ , then the configuration  $s'$  obtained by restricting  $\mathbf{a}(a) = s$  to the  $\Sigma_I$ -substructure generated by the  $\mathbf{a}(\underline{i})$  is equivalent to  $s$  by Lemma 5.8 and hence belongs to a basis of  $K$ . Thus Lemma 5.11 applies and gives  $(\mathcal{M}, \mathbf{a}) \models Min(\phi, a, \underline{i})$ .  $\square$

The next step towards the goal of obtaining a completely symbolic method for mechanizing the result stated in Theorem 5.10 consists of finding a purely symbolic substitute of the function  $\gamma$  used in the proof of Theorem 5.10. The following result is the key to achieve this.

**Proposition 5.14.** *Let  $T_E$  be locally finite,  $K := \exists \underline{i} \phi(\underline{i}, a[\underline{i}])$  be an  $\exists^I$ -formula, and  $L$  be an  $\exists^I$ -formula. The following two conditions are equivalent:*

- (i) *for every  $s$  in a basis for  $K$ , there exists a configuration  $s'$  in a basis for  $L$  such that  $s \leq s'$ ;*
- (ii)  *$L$  is (up to  $A_I^E$ -equivalence) of the form  $\exists \underline{i}, \underline{j} \psi(\underline{i}, \underline{j}, a[\underline{i}], a[\underline{j}])$  for a quantifier-free formula  $\psi$  and*

$$\text{if } A_I^E \models Min(\psi, a, \underline{i}, \underline{j}) \rightarrow \theta(\underline{t}, a[\underline{t}]) \quad \text{then} \quad A_I^E \models Min(\phi, a, \underline{i}) \rightarrow \theta(\underline{t}, a[\underline{t}]),$$

*for all quantifier free  $(\Sigma_E \cup \Sigma_I)$ -formula  $\theta$  and for all tuple of terms  $\underline{t}$  taken from the set of the representative  $\Sigma_I(\underline{i})$ -terms.*

*Proof.* Assume (i). We first apply a syntactic transformation to  $L$  as follows. Let  $B, B'$  be bases for  $K, L$ , respectively; we know that for every  $(s, \mathcal{M}_s) \in B$  there is  $(s^L, \mathcal{M}_s^L) \in B'$  such that  $s \leq s^L$ : the relationship  $s \leq s^L$  is due to the existence of a pair of embeddings  $(\mu_s, \nu_s)$  as required by the configuration ordering definition. For every  $s \in B$  and for every assignment  $\mathbf{a}$  such that  $\mathbf{a}(a) = s$  and  $(\mathcal{M}_s, \mathbf{a}) \models \phi(\underline{i}, a[\underline{i}])$ , we build the diagram formula  $K_{\mathbf{a}}$  for  $s^L$  given by

$$\exists \underline{i} \exists \underline{k} (\delta_{s^L}(\underline{i}, \underline{k}) \wedge \delta_{s^E}(a[\underline{i}], a[\underline{k}])) \tag{5.8}$$

where the variables  $\underline{k}$  are names for the elements in the complement subset  $supp(s^L) \setminus \mu_s(\mathbf{a}(\underline{i}))$  (here  $supp(s^L)$  is the support of the  $\Sigma_I$ -structure  $s^L$ ). Notice that the formula (5.8) is nothing but formula (4.2) used in the proof of Proposition 4.5(i).<sup>13</sup> Since, for a configuration  $t$ , the fact that  $t \in \llbracket K_{\mathbf{a}} \rrbracket$  means that there are suitable embeddings witnessing that  $s^L \leq t$ , we have that  $\llbracket L \rrbracket = \llbracket L \vee \bigvee_{\mathbf{a}} K_{\mathbf{a}} \rrbracket$ , hence by Proposition 4.2 the formula  $L$  is  $A_I^E$ -equivalent to

<sup>13</sup>It might happen here that duplicate variables are used because the  $\mathbf{a}(\underline{i})$  need not be distinct. This is not a problem: if different index variables (say  $i_1, i_2$ ) naming the same element are employed, the diagram formula will contain a conjunct like  $i_1 = i_2$ . The embedding property of Robinson Diagram Lemma is not affected by these duplications.

$L \vee \bigvee_{\mathbf{a}} K_{\mathbf{a}}$ .<sup>14</sup> Up to logical equivalence, we can move the existentially quantified variables outside the disjunctions so that  $L$  is equivalent to a prenex existential formula of the kind  $\exists \underline{i} \exists \underline{j} \psi$ . With this new syntactic form, the following property holds: for every  $s \in B$  and for every assignment  $\mathbf{a}$  such that  $\mathbf{a}(a) = s$  and  $(\mathcal{M}_s, \mathbf{a}) \models \phi(\underline{i}, a[\underline{i}])$ , there is an assignment  $\mathbf{a}^L$  such that (i)  $(\mathcal{M}_s^L, \mathbf{a}^L) \models \psi(\underline{i}, \underline{j}, a[\underline{i}], a[\underline{j}])$ , (ii)  $\mathbf{a}^L(\underline{i}) = \mu_s(\mathbf{a}(\underline{i}))$ , and (iii)  $\mathbf{a}^L(a) = s^L$ . Since  $s^L$  is in a basis of  $L$ , from Corollary 5.13, it follows also that  $(\mathcal{M}_s^L, \mathbf{a}^L) \models \text{Min}(\psi, a, \underline{i}, \underline{j})$ .

Suppose now that  $A_I^E \not\models \text{Min}(\phi, a, \underline{i}) \rightarrow \theta(\underline{t}(\underline{i}), a[\underline{t}(\underline{i})])$ ; by Lemma 5.11 (and by the fact that  $\phi, \theta$  are quantifier-free) this means that there are a configuration  $(s, \mathcal{M}_s) \in B$  and an assignment  $\mathbf{a}$  such that  $(\mathcal{M}_s, \mathbf{a}) \models \phi(\underline{i}, a[\underline{i}])$  and  $(\mathcal{M}_s, \mathbf{a}) \not\models \theta(\underline{t}, a[\underline{t}])$ . Since  $\theta$  is quantifier-free, taking the assignment  $\mathbf{a}^L$  satisfying (i)-(ii)-(iii) above, we get that  $(\mathcal{M}_s^L, \mathbf{a}^L) \not\models \theta(\underline{t}, a[\underline{t}])$ , thus also  $(\mathcal{M}_s^L, \mathbf{a}^L) \not\models \text{Min}(\psi, a, \underline{i}, \underline{j}) \rightarrow \theta(\underline{t}, a[\underline{t}])$ .

Conversely, assume (ii). Fix  $(s, \mathcal{M}_s)$  in a basis  $B$  for  $K$  and an assignment  $\mathbf{a}$  such that  $(\mathcal{M}_s, \mathbf{a}) \models \phi(\underline{i}, a[\underline{i}])$ ; by Corollary 5.13, we have that  $(\mathcal{M}_s, \mathbf{a}) \models \text{Min}(\phi, a, \underline{i})$ . Let  $\underline{t}$  be the representative  $\Sigma_I(\underline{i})$ -terms and let  $\theta(\underline{t}(\underline{i}), a[\underline{t}(\underline{i})])$  be the negation of the formula  $\delta_{s_I}(\underline{t}(\underline{i})) \wedge \delta_{s_E}(a[\underline{t}(\underline{i})])$ . We have  $(\mathcal{M}_s, \mathbf{a}) \not\models \text{Min}(\phi, a, \underline{i}) \rightarrow \theta(\underline{t}, a[\underline{t}])$ , hence there are  $\mathcal{N}$  and  $\mathbf{b}$  such that  $(\mathcal{N}, \mathbf{b}) \not\models \text{Min}(\psi, a, \underline{i}, \underline{j}) \rightarrow \theta(\underline{t}, a[\underline{t}])$ . By restricting the support of  $\mathcal{N}_I$  if needed, we can suppose that  $\mathcal{N}$  is a finite index model and that  $\mathcal{N}_I$  is generated by the elements assigned by  $\mathbf{b}$  to the  $\underline{i}, \underline{j}$ . Let  $s'$  be  $\mathbf{b}(a)$ : from Lemma 5.11 it follows that  $s'$  is in a basis for  $L$ ; also, from the fact that  $(\mathcal{N}, \mathbf{b}) \not\models \theta(\underline{t}, a[\underline{t}])$ , we can conclude that  $s \leq s'$ , as desired.  $\square$

In the following, we will write  $K \leq L$  whenever one of the (equivalent) conditions in Proposition 5.14 holds. We show that, under the working assumption that  $T_E$  is locally finite, it is possible to compute all the finitely many (up to  $A_I^E$ -equivalence)  $\exists^I$ -formulae  $K$  such that  $K \leq L$ .

**Proposition 5.15.** *Let  $T_E$  be locally finite. Given an  $\exists^I$ -formula  $L$ , there are only finitely many (up to  $A_I^E$ -equivalence)  $\exists^I$ -formulae  $K$  such that  $K \leq L$  and all such  $K$  can be effectively computed.*

*Proof.* Suppose that  $L$  is of the form  $\exists \underline{k} \gamma$ . To use the criterion of Proposition 5.14(ii) in an effective way, we only need to find a bound for the length of the tuples  $\underline{i}$  and  $\underline{j}$ . In fact, once the bound is known the search space for formulae of the forms  $\exists \underline{i} \exists \underline{j} \psi$  and  $\exists \underline{i} \phi$  satisfying the conditions (which can be effectively checked by using Theorem 3.3)

$$\begin{aligned} A_I^E \models \exists \underline{k} \gamma &\leftrightarrow \exists \underline{i} \exists \underline{j} \psi, & \text{and for all } \theta(\underline{t}(\underline{i}), a[\underline{t}(\underline{i})]) \\ A_I^E \models \text{Min}(\psi, a, \underline{i}, \underline{j}) \rightarrow \theta(\underline{t}, a[\underline{t}]) &\Rightarrow A_I^E \models \text{Min}(\phi, a, \underline{i}) \rightarrow \theta(\underline{t}, a[\underline{t}]) \end{aligned}$$

is finite. This is because  $T_I$  and  $T_E$  are both locally finite and hence, there are only finitely many quantifier-free formulae of the required type involving a fixed number of index variables which are not  $A_I^E$ -equivalent. The proof of Proposition 5.14 shows that the lengths of  $\underline{i}$  and  $\underline{j}$  are both bounded by the maximum cardinality  $N$  of the support of  $s_I$ , where  $s_I$  is a configuration that belongs to a basis for  $L \equiv \exists \underline{k} \gamma$ . For  $\underline{j}$ , this is clear from the proof itself while for  $\underline{i}$ , it is a consequence of the following considerations. First, we can restrict the search to formulae  $K$  of the form  $\exists \underline{i} \phi$ , where the length of  $\underline{i}$  is minimal, i.e.  $K$  is not be equivalent to a formula with a shorter existential prefix. Furthermore, by Proposition 4.5,  $K$  is equivalent to  $K_{s_1} \vee \dots \vee K_{s_n}$ , where  $\{s_1, \dots, s_n\}$  is a basis for  $K$ . In turn, by (4.2), this means that there must exist a configuration  $t$  in a basis for  $K$  such that the cardinality

<sup>14</sup>The assignments are infinite, but only finitely many variables are mentioned in them, so that only finitely many formulae  $K_{\mathbf{a}}$  can be produced.

of  $t_I$  is bigger than or equal to the length of  $\underline{i}$ ; since  $t \leq s$  for some  $s$  in a basis for  $L$  (see Proposition 5.14(ii)), we have that the length of  $\underline{i}$  cannot exceed  $N$ . To conclude, it is sufficient to observe that  $N$  cannot be bigger than the number of the representative  $\Sigma_I(\underline{k})$ -terms.  $\square$

**Definition 5.16.** We say that  $K$  covers  $L$  iff both  $K \leq L$  and  $A_I^E \models L \rightarrow K$ .

The following example illustrates the notions just introduced and will be useful also when discussing the implementation of our invariant synthesis technique (see Section 6.2 below).

**Example 5.17.** Let  $\Sigma_I$  be relational and  $T_E$  be a locally finite theory admitting elimination of quantifiers. Let

$$L := \exists \underline{i}, \underline{j}. (\psi_E(a[\underline{i}], a[\underline{j}]) \wedge \psi_I(\underline{i}, \underline{j}) \wedge \delta_I(\underline{i})) \quad (5.9)$$

be a primitive differentiated and  $A_I^E$ -satisfiable  $\exists^I$ -formula such that (i)  $\underline{i} \cap \underline{j} = \emptyset$ , (ii)  $\psi_E(\underline{e}, \underline{d})$  is a conjunction of  $\Sigma_E$ -literals; (iii)  $\psi_I(\underline{i}, \underline{j})$  is a conjunction of  $\Sigma_I$ -literals; (iv)  $\delta_I(\underline{i})$  is a maximal conjunction of  $\Sigma_I(\underline{i})$ -literals (i.e. for every  $\Sigma(\underline{i})$ -atom  $A(\underline{i})$ ,  $\delta_I$  contains either  $A(\underline{i})$  or its negation). If

$$K := \exists \underline{i} (\delta_I(\underline{i}) \wedge \phi_E(a[\underline{i}])), \quad (5.10)$$

where  $\phi_E(\underline{e})$  is  $T_E$ -equivalent to  $\exists \underline{d} \psi_E(\underline{e}, \underline{d})$ ,<sup>15</sup> then  $K$  covers  $L$  and in particular  $K \leq L$ . We prove this fact in the following.

*Proof.* We use Proposition 5.14(ii): as shown in Remark 5.12, since  $L$  and  $K$  are differentiated, we can avoid mentioning the corresponding formulae  $Min$  in the condition of Proposition 5.14(ii) and just prove that

$$\begin{aligned} A_I^E \not\models \delta_I(\underline{i}) \wedge \phi_E(a[\underline{i}]) \rightarrow \theta(\underline{i}, a[\underline{i}]) &\Rightarrow \\ A_I^E \not\models \delta_I(\underline{i}) \wedge \psi_I(\underline{i}, \underline{j}) \wedge \psi_E(a[\underline{i}], a[\underline{j}]) \rightarrow \theta(\underline{i}, a[\underline{i}]) & \end{aligned}$$

for every  $\theta$  (notice that, since  $\Sigma_I$  is relational, the only  $\Sigma_I(\underline{i})$ -terms are the  $\underline{i}$ ). Pick a model  $\mathcal{M}$  and an assignment  $\mathbf{a}$  such that  $(\mathcal{M}, \mathbf{a}) \models \delta_I(\underline{i}) \wedge \phi_E(a[\underline{i}])$  and  $(\mathcal{M}, \mathbf{a}) \not\models \theta(\underline{i}, a[\underline{i}])$ . We can freely assume that the support of  $\mathcal{M}_I$  is a  $\Sigma_I$ -structure generated by the  $\mathbf{a}(\underline{i})$ ; by modifying the value of  $\mathbf{a}$  on the element variables  $\underline{d}$ , if needed, we can also assume that  $(\mathcal{M}, \mathbf{a}) \models \psi_E(a[\underline{i}], \underline{d})$  (this is because  $\phi_E(\underline{e})$  is  $T_E$ -equivalent to  $\exists \underline{d} \psi_E(\underline{e}, \underline{d})$ ). Since  $L$  is consistent, there are also a model  $\mathcal{N}$  and an assignment  $\mathbf{b}$  such that  $(\mathcal{N}, \mathbf{b}) \models \delta_I(\underline{i}) \wedge \psi_I(\underline{i}, \underline{j}) \wedge \psi_E(a[\underline{i}], a[\underline{j}])$ . Again, we can assume that the support of  $\mathcal{N}_I$  is a  $\Sigma_I$ -structure generated by the  $\mathbf{a}(\underline{i}, \underline{j})$ ; since  $\delta_I(\underline{i})$  is maximal, it is a diagram formula, hence (up to an isomorphism)  $\mathcal{M}_I$  is a substructure of  $\mathcal{N}_I$ . Let us now take the model  $\mathcal{N}'$ , whose  $\Sigma_I$ -reduct is  $\mathcal{N}_I$  and whose  $\Sigma_E$ -reduct is  $\mathcal{M}_E$ . Let  $\mathbf{b}'$  be the assignment which is like  $\mathbf{b}$  as far as the index variables  $\underline{i}, \underline{j}$  are concerned and which associates with the variable  $a$  the array whose  $\mathbf{b}'(\underline{i})$ -values are the  $\mathbf{b}'(\underline{i}) = \mathbf{a}(\underline{i})$ -values of  $\mathbf{a}(a)$  and whose  $\mathbf{b}'(\underline{j})$ -values are the  $\underline{d}$  (notice that this is correct because by differentiatedness of  $L$  the  $\mathbf{b}'(\underline{ij})$  are all distinct). It turns out that  $(\mathcal{N}', \mathbf{b}') \not\models \delta_I(\underline{i}) \wedge \psi_I(\underline{i}, \underline{j}) \wedge \psi_E(a[\underline{i}], a[\underline{j}]) \rightarrow \theta(\underline{i}, a[\underline{i}])$ , as desired.  $\square$

<sup>15</sup> $\phi_E$  is guaranteed to exist as  $T_E$  admits elimination of quantifiers.

We are now in the position to take the final step towards the goal of obtaining a completely symbolic method for restating the results from Theorem 5.10. Let  $\text{ChooseCover}(L)$  be a procedure that returns non-deterministically one of the  $\exists^I$ -formulae  $K$  such that  $K$  covers  $L$  (this procedure is playing the role of the function  $\gamma$  from the proof of Theorem 5.10). We consider the procedure  $\text{SInv}$  in Figure 1 (b) for the computation of safety invariants and prove its correctness.

**Theorem 5.18.** *Let  $T_E$  be locally finite. Then, there exists a safety invariant for  $U$  iff the procedure  $\text{SInv}$  in Figure 1 (b) returns a safety invariant for  $U$ , for a suitable  $\text{ChooseCover}$  function.*

*Proof.* Suppose that  $\text{SInv}$  returns  $B$  after  $k + 1$  iterations of the loop: we show that  $\neg B$  is a safety invariant. Notice that  $B$  is a disjunction  $P_0 \vee \dots \vee P_k$  of  $\exists^I$ -formulae such that for all  $i = 0, \dots, k$ ,

(I): the formula  $I \wedge P_i$  is not  $A_I^E$ -satisfiable;

also  $P_i$  covers  $\text{Pre}(\tau, P_{i-1})$  and  $P_0$  covers  $U$ , which means in particular that

(II):  $A_I^E \models \forall a (\text{Pre}(\tau, P_{i-1})(a) \rightarrow P_i(a))$  and  $A_I^E \models \forall a (U(a) \rightarrow P_0(a))$ .

Finally,  $\text{SInv}$  could exit the loop because for some  $P_{k+1}$  covering  $\text{Pre}(\tau, P_k)$ , it happened that  $P_{k+1} \wedge \neg B$  was not  $A_I^E$ -satisfiable: these two conditions entail that

(III):  $A_I^E \models \forall a (\text{Pre}(\tau, P_k)(a) \rightarrow B(a))$ .

Conditions (i) and (iii) of Definition 5.1 now easily follows from (I) and (II); we only need to check condition (ii) of Definition 5.1, namely (up to logical equivalence) that  $A_I^E \models \forall a (\text{Pre}(\tau, B)(a) \rightarrow B(a))$ : since  $\text{Pre}(\tau, B)$  is logically equivalent to the disjunction  $\bigvee_{i=0}^n \text{Pre}(\tau, P_i)$ , the claim follows immediately from (II)-(III).

Let us now prove the converse, i.e. that in case a safety invariants exists,  $\text{SInv}$  is able to compute one. Recall the proof of Theorem 5.10: given the negation  $H$  of a safety invariant for  $U$ , another negation  $L$  of a safety invariant for  $U$  is produced in the following way. Define the sequence of  $\exists^I$ -formulae  $L_i$  as follows: (i)  $L_0 := \gamma(U)$  and (ii)  $L_{i+1} := L_i \vee \gamma(\text{Pre}(\tau, L_i))$ . Our  $L$  is the  $L_i$  with the smallest  $i$  such that  $L_{i+1}$  is  $A_I^E$ -equivalent to  $L_i$  (the proof of Theorem 5.10 guarantees that such an  $i$  exists).

The above recursive definition for  $L_i$  is based on the function  $\gamma$ , which is defined (non symbolically) by making use of configurations. Actually, for an  $\exists^I$ -formula  $S$  such that  $\llbracket S \rrbracket \subseteq \llbracket H \rrbracket$ , the function  $\gamma(S)$  returns an  $\exists^I$ -formula  $K_{a_1} \vee \dots \vee K_{a_n}$ , where  $\{a_1, \dots, a_n\} \subseteq \llbracket H \rrbracket$  is a minimal set of configurations taken from a basis of  $H$  such that  $\llbracket S \rrbracket \subseteq \uparrow a_1 \cup \dots \cup \uparrow a_n$ . Using Proposition 5.14, it is not difficult to see that minimality implies  $\gamma(S) \leq S$ : in fact, condition  $\llbracket S \rrbracket \subseteq \uparrow a_1 \cup \dots \cup \uparrow a_n$  says that for every  $s$  in a basis for  $S$  there is  $a_i$  in the basis  $\{a_1, \dots, a_n\}$  for  $\gamma(S)$  such that  $a_i \leq s$ , but the converse (which is what really matters for us in view of Proposition 5.14(i)) must hold too, by minimality. This can be shown as follows: if any  $a_i$  is eliminated, the relation

$$\llbracket S \rrbracket \subseteq \uparrow a_1 \cup \dots \cup \uparrow a_{i-1} \cup \uparrow a_{i+1} \cup \dots \cup \uparrow a_n$$

does no longer hold, hence there is an  $s$  from a basis of  $S$  such that  $a_j \not\leq s$  for all  $j = 1, \dots, i-1, i+1, \dots, n$ . Since, on the contrary,  $\llbracket S \rrbracket \subseteq \uparrow a_1 \cup \dots \cup \uparrow a_n$  holds, we must conclude that  $a_i \leq s$ . Hence for every  $a_i$  there is an  $s$  in a basis of  $S$  such that  $a_i \leq s$ .

Thus  $\gamma(S)$  is such that  $\gamma(S) \leq S$  and  $A_I^E \models S \rightarrow \gamma(S)$ , i.e.  $\gamma(S)$  covers  $S$ . It is then clear that an appropriate choice of the function  $\text{ChooseCover}$  in  $\text{SInv}$  can return precisely

the formulae  $L_i$  so that they are assigned to the variable  $B$  at the  $i$ th-loop of the procedure, thus justifying the claim of the Theorem.  $\square$

When  $\text{ChooseCover}(L) = L$ , i.e.  $\text{ChooseCover}$  is the identity (indeed,  $L$  covers  $L$ ), the procedure  $\text{SInv}$  is the (exact) dual of  $\text{BReach}$  in Figure 1 (a) and, hence it can only return (the negation of) a symbolic representation of all backward reachable states as a safety invariant.

## 6. PRAGMATICS OF INVARIANT SYNTHESIS AND EXPERIMENTS

The main drawback of algorithm  $\text{SInv}$  (in Figure 1 (b), explained in the last section) is the non determinism of the function  $\text{ChooseCover}$ . Although finite, the number of formulae covering a certain  $\exists^I$ -formula is so large to make any concrete implementation of  $\text{SInv}$  impractical. Instead, we prefer to study how to integrate the synthesis of invariants into the backward reachability algorithm of Figure 1 (a). Given that finding a *safety invariant* could be infeasible through an exhaustive search, we content ourselves to find invariants *tout court* and use them to prune the search space of the backward reachability algorithm  $\text{BReach}$  (in Figure 1 (a)).

**6.1. Integrating Invariant Synthesis within Backward Reachability.** In our symbolic framework, at the  $n$ -th iteration of the loop of the procedure  $\text{BReach}$ , the set of backward reachable states is represented by the formula stored in the variable  $B$  (which is equivalent to  $BR^n(\tau, U)$ ). So, ‘pruning the search space of the backward reachability algorithm’ amounts to disjoining the negation of the available invariants to  $B$ . In this way, the extra information encoded in the invariants makes the satisfiability test at line 2 (for fix-point checking) more likely to be successful and possibly decreases the number of iterations of the loop. Indeed, the problem is to synthesize such invariants. Let us consider this problem at a very abstract level.

Suppose the availability of a function  $\text{Choose}$  that takes an  $\exists^I$ -formula  $P$  and returns a (possibly empty) finite set  $S$  of  $\exists^I$ -formulae representing ‘useful (with respect to  $P$ ) candidate invariants.’ We can integrate the synthesis of invariants within the backward reachability algorithm by adding between lines 4 and 5 in Figure 1 (a) the following instructions:

```
4'      foreach  $CINV \in \text{Choose}(P)$  do
          if  $\text{BReach}(CINV) = (\text{safe}, B_{CINV})$  then  $B \leftarrow B \vee \neg B_{CINV}$ ;
```

where  $CINV$  stands for ‘candidate invariant.’ The resulting procedure will be indicated with  $\text{BReach+Inv}$  in the following. Notice that  $\text{BReach}$  is used here as a sub-procedure of  $\text{BReach+Inv}$ .

**Proposition 6.1.** *If the procedure  $\text{BReach+Inv}$  terminates by returning safe (unsafe), then  $S$  is safe (unsafe) with respect to  $U$ .*

*Proof.* The claim is trivial when  $\text{BReach+Inv}$  returns *unsafe*. Let us consider the situation when the procedure terminates by returning *safe* at the  $(k + 1)$ -th iteration of the main loop. Observe that the content of the variable  $B$  is

$$Pre^0(\tau, U) \vee Pre^1(\tau, U) \vee \dots \vee Pre^k(\tau, U) \vee H_1 \vee \dots \vee H_m \quad (6.1)$$



at the  $(k + 1)$ -th iteration of the loop, where  $H_1, \dots, H_m$  are negations of invariants (see Property 5.3). For reductio, suppose that the system is unsafe, i.e. for some  $n \geq 0$ , the formula (3.2) (shown here for the sake of readability)

$$I(a_n) \wedge \tau(a_n, a_{n-1}) \wedge \dots \wedge \tau(a_1, a_0) \wedge U(a_0)$$

is  $A_I^E$ -satisfiable. Assume that the formula is true in a model of  $A_I^E$  with the array assignments  $s_n, \dots, s_0$ ; in the following, we say that  $s_n, \dots, s_0$  is a *bad trace*. We also assume that  $s_n, \dots, s_0$  is a bad trace of shortest length. Since the formulae  $I \wedge Pre^0(\tau, U)$ ,  $I \wedge Pre^1(\tau, U)$ ,  $\dots$ , and  $I \wedge Pre^k(\tau, U)$  are all  $A_I^E$ -unsatisfiable (see line 3 of Figure 1 (a), which is also part of BReach+Inv), it must be  $n > k$ . Let us now focus on  $s_{k+1}$ ; since BReach+Inv returned safe at iteration  $k + 1$ , it must have exited the loop because the formula currently stored in  $P$  (which is  $Pre^{k+1}(\tau, U)$ ) is not  $A_I^E$ -satisfiable with the negation of the formula currently in  $B$  (which is (6.1)). Hence,  $s_{k+1}$  (which satisfies  $Pre^{k+1}(\tau, U)$ ) must satisfy either some  $Pre^l$  (for  $l < k + 1$ ) or some  $H_i$ , but both alternatives are impossible. In fact, the former would yield a shorter bad trace, whereas the latter is in contrast to the fact that  $s_{k+1}$  is forward reachable from a state satisfying  $I$  and, as such, it should satisfy the invariant  $\neg H_i$ .  $\square$

The procedure BReach+Inv is

- incomplete, in the sense that it is not guaranteed to terminate even when a safety invariant exists,
- deterministic, since no backtracking is required,
- highly parallelizable: it is possible to run in parallel as many instances of BReach as formulae in the set returned by Choose, and
- it performs well (for *appropriate* Choose functions, see below for a discussion of the meaning of “appropriate” in this context) as witnessed by the experiments in the next section.

As a result, invariant synthesis becomes a powerful *heuristic* within a refined version of the basic backward reachability algorithm. Furthermore, its integration in the tableaux calculus of Section 3.3 is particularly easy: *just use the calculus itself with some bounds on the resources* (such as a limit on the depth of the tree) to check if a candidate invariant is a “real” invariant. Indeed, the crucial point is how to design an *appropriate* function Choose. There are several possible criteria leading to a variety of implementations for Choose. The usefulness of the resulting functions is likely to depend on the application. Despite the complexity of the design space, it is possible to identify a *minimal requirement* on Choose by taking into account the tableaux calculus introduced in Section 3.3. To this end, recall that backward reachable sets of states are described by primitive differentiated formulae and that a formula  $P$  representing a pre-image is eagerly expanded to disjunctions of primitive differentiated formulae by using the Beta rule. Thus, a reasonable implementation of Choose should be such that  $\text{Choose}(P) = S$  where  $S$  is a set of primitive differentiated formulae such that each  $Q' \in S$  is implied by a disjunct  $Q$  occurring in the disjunction of primitive differentiated formulae obtained as expansion of  $P$ . In this way, each  $Q' \in S$  can be seen as a *tentative over-approximation* of  $Q$ . (Notice that guessing a candidate invariant can be seen as a form of abstraction.) All the implementations of the function Choose in MCMT satisfy the minimal requirement above and can be selected by appropriate command line options and directives to be included in the input file (the interested reader is pointed to the user manual available in the distribution for details). We now describe two types of

abstractions that lead to different implementations of the function `Choose` that are available in the current release of MCMT.

**6.2. Index Abstraction.** Index abstraction amounts to eliminating some index variables; if done in the appropriate way, this is equivalent to replacing configurations with sub-configurations (as discussed in Section 5). Thus, it is possible to design approximations (quite loose, but suitable for implementation) of the procedure suggested in the proof of Theorem 5.18. An idea (close to what is implemented in the current release of MCMT) is to follow the suggestions in Example 5.17 so as to satisfy the minimal requirement discussed above on `Choose`. More precisely, given  $Q := \exists \underline{k}.\theta(\underline{k}, a[\underline{k}])$ , we first try to transform it into the form of (5.9), i.e.

$$\exists \underline{i} \underline{j} . (\psi_E(a[\underline{i}], a[\underline{j}]) \wedge \psi_I(\underline{i}, \underline{j}) \wedge \delta_I(\underline{i})).$$

To do this, we decompose  $\underline{k}$  into two disjoint sub-sequences  $\underline{i}$  and  $\underline{j}$  such that  $\underline{k} = \underline{i} \cup \underline{j}$  according to some criteria: if the conjunction of  $\Sigma_I(\underline{i})$  literals occurring in  $\theta$  is maximal, we get a candidate invariant by returning the corresponding  $\exists^I$ -formula (5.10), i.e.

$$\exists \underline{i} (\delta_I(\underline{i}) \wedge \phi_E(a[\underline{i}])).$$

This is computationally feasible in many situations. For example, quantifier elimination reduces to a trivial substitution if  $T_E$  is an enumerated data-type theory and the  $\Sigma_E$ -literals in  $\theta$  (i.e. those in  $\psi_E$ ) are all positive. The maximality of  $\theta$  is guaranteed (by being differentiated) if  $T_I$  is the theory of finite sets. Another case in which maximality of  $\theta$  is guaranteed is when  $T_I$  is the theory of linear orders and  $\underline{i} = i_1$  or ( $\underline{i} = i_1, i_2$  and  $\theta$  contains the atom  $i_1 < i_2$ ). In more complex cases, it is possible to obtain a useful formula (similar to (5.10)) in a purely syntactic and computationally cheap way. There is no risk in using methods giving very coarse approximations since a candidate invariant is used for pruning the search space of the backward reachability procedure only if it has been proved to be a “real” invariant (see also Remark 6.2 below).

**6.3. Signature Abstraction.** Index abstraction can be useless or computationally too expensive (if done precisely) in several applications. Even worse, when  $T_E$  is not locally finite, the related notion of sub-configuration loses most of its relevance. In these cases, other forms of abstraction inspired to predicate abstraction [44] may be of great help. Although predicate abstraction with refinement (as in the CEGAR loop) is not yet implemented in MCMT, it features a technique for invariant synthesis that we have called *signature abstraction*, which can be seen as a simplified version of predicate abstraction. This technique uses quantifier elimination (whenever possible) to eliminate the literals containing a selected sub-set  $X$  of the set of array variables. The subset  $X$  can either be suggested by the user or dynamically built by the tool from the shape of the disjunct belonging to the pre-image being currently computed. Again, the elimination is applied to each of the primitive differentiated disjuncts of the currently computed pre-image  $P$  to obtain the differentiated formulae to form the set of formulae returned by `Choose`. It is easy to see that this way of implementing the function `Choose` satisfies the minimal requirement discussed above.

**Remark 6.2.** The reader may wonder whether the use of abstraction techniques can have a negative impact on the correctness of MCMT outcome. We emphasize that this is not the case because of the way the candidate invariants are used to prune the search space during backward reachability. In fact, abstraction is just to generate the candidate invariants

which are then tested to be “real” invariants by a resource bounded version of backward reachability. Only if candidate invariants pass this test, they are used to prune the search of backward reachable states. In other words, the answer supplied by MCMT to a safety problem is always correct: as it is clear from the proof of Proposition 6.1, the set of backward reachable states can be augmented if invariants are used during backward search, but it is augmented by adding it only states satisfying the negation of an invariant (these states are not forward reachable, hence they cannot alter safety checks). As a consequence, safety tests remain exhaustive, although it may happen that resources (such as computation time) are wasted in checking candidate invariants that turn out not to be “real” invariants or not to be useful to significantly prune the search space.

**6.4. Experiments.** To show the flexibility and the performances of MCMT, we have built a library of benchmarks in the format accepted by our tool by translating from a variety of sources safety problems. More precisely, our sources were the following:

- parametrised systems from the distribution of the infinite model checkers PFS (<http://www.it.uu.se/research/docs/fm/apv/tools/pfs>) and Undip (<http://www.it.uu.se/research/docs/fm/apv/tools/undip>),
- parametrised and distributed systems from the invisible invariant methods (see, e.g., [12]),
- imperative programs manipulating arrays (such as sorting or string manipulation) taken from standard books about algorithms,
- imperative programs manipulating numeric variables from the distribution of the model checker ARMC (<http://www7.in.tum.de/~rybal/armc>),
- protocols from the distribution of Mur $\phi$  extended with predicate abstraction (<http://verify.stanford.edu/satyaki/research/PredicateAbstractionExamples.html>).

We did not try to be exhaustive in the selection of problems but rather to pick problems from the wider possible range of different classes of infinite state systems so as to substantiate the claim about the flexibility of our tool. All the files in MCMT format are contained in the MCMT distribution which is available at the tool web page (<http://homes.dsi.unimi.it/~ghilardi/mcmt>). Each file comes with the indication of source from which it has been adapted and a brief informal explanation about its content.

We divided the problems into four categories: mutual exclusion and cache coherence protocols taken mainly from the distributions of PFS and Undip (see Tables 1 and 2), imperative programs manipulating arrays (see Table 3), and heterogeneous problems (see Table 4) taken from the remaining sources listed above. For the first two categories, the benchmark set is sufficiently representative, whereas for the last two categories just some interesting examples have been submitted to the tool. For each category, we tried the tool in two configurations: one, called “Default Setting,” is the standard setting used when MCMT is invoked without any option and the other, called “Best Setting,” is the result of some experimentation with various heuristics for invariant synthesis, signature abstraction, and acceleration. It is possible that for some problems, the “real” best setting is still to be identified and the results reported here can be further improved.

In Tables 1, 2, 3, and 4, the column ‘d’ is the depth of the tableaux obtained by applying the rules listed in Section 3.3, ‘#n’ is the number of nodes in the tableaux, ‘#d’ is the number of nodes which are deleted because they are subsumed by the information contained in the others (see [40] for details about this point), ‘#SMT’ is the number of invocations to Yices

Table 1: Mutual exclusion protocols

Problem	Default Setting					Best Setting					
	d	#n	#d	#SMT	time	d	#n	#d	#SMT	#inv.	time
Bakery	2	1	0	6	0.00	2	1	0	6	0	0.00
Bakery_bogus	8	90	14	1413	0.81	8	53	4	1400	7	0.68
Bakery_e	12	48	17	439	0.20	7	8	1	213	16	0.10
Bakery_Lamport	12	56	15	595	0.27	4	7	1	209	7	0.08
Bakery_t	9	28	5	251	0.11	7	8	1	134	5	0.06
Burns	14	56	7	373	0.14	2	2	1	53	3	0.02
Dijkstra	14	122	37	2920	2.11	2	1	1	215	12	0.08
Dijkstra1	13	38	11	222	0.10	2	1	1	35	2	0.02
Distrib_Lamport	23	913	242	47574	120.62	23	248	42	19254	7	32.84
Java M-lock	9	23	2	289	0.10	9	23	2	289	0	0.10
Mux_Sem	7	8	2	57	0.02	2	1	1	65	6	0.02
Rickart_Agrawala	13	458	119	35355	187.04	13	458	119	35355	0	187.04
Sz_fp	22	277	3	7703	5.12	22	277	3	7703	0	5.12
Sz_fp_ver	30	284	38	10611	6.66	30	284	38	10611	0	6.66
Szymanski	17	136	10	2529	1.60	9	14	5	882	12	0.30
Szymanski_at	23	1745	311	424630	540.19	9	22	10	2987	42	1.25
Ticket	9	18	0	284	0.17	9	18	0	284	0	0.17

during backward reachability to solve fix-point and safety checks, ‘#inv.’ is the number of invariants found by the available invariant synthesis techniques (see also [39] for a more in-depth discussion on some of these issues),<sup>16</sup> and ‘time’ is the total amount of time (in seconds) taken by the tool to solve the safety problem. Timings were obtained on a Intel Centrino 1.729 GHz with 1 Gbyte of RAM running Linux Gentoo. In some cases, the system seemed to diverge as it clearly entered in a loop: it kept applying the same sequence of transitions. In these cases, we stopped the system, left the corresponding line of the table empty, and put ‘timeout’ in the last column.

As it is apparent by taking a look at the Tables, gaining some expertise in using the available options of the tool may give dramatic improvements in performances, either in terms of reduced timings or in getting the system to terminate. For the category “Mutual exclusion protocols,” invariant synthesis is helpful to reduce the solving time for the larger examples. For the category “Cache coherence protocols,” the effect of invariant synthesis as well as other techniques is negligible. For the category “Imperative programs,” invariant synthesis techniques are the key to make the tool terminate on almost all problems. In particular, signature abstraction, introduced in this last version of the tool, is a crucial ingredient.

A comparative analysis is somewhat difficult in lack of a standard for the specifications of safety problems. This situation is similar to the experimental evaluation of SMT solvers before the introduction of the SMT-LIB standard [53]. It would be interesting to investigate if the proposed format can become the new interlingua for infinite state model checkers so that exchange of problems becomes possible as well as the fair comparison of performances. Just to give an idea of the relative performance of our tool, we only mention that MCMT

<sup>16</sup>In the table for the “Default Setting,” the column labelled with ‘#inv.’ is not present because MCMT’s default is to turn off invariant synthesis.

Table 2: Cache coherence protocols

Problem	Default Setting					Best Setting					
	d	#n	#d	#SMT	time	d	#n	#d	#SMT	#inv.	time
Berkeley	2	1	0	16	0.00	2	1	0	16	0	0.00
Futurebus	8	37	3	998	0.96	8	37	3	998	0	0.96
German07	26	2442	576	121388	145.68	26	2442	576	121388	0	145.68
German_buggy	16	1631	203	41497	49.70	16	1631	203	41497	0	49.70
German_ca	9	13	0	62	0.03	9	13	0	62	0	0.03
German_pfs	33	11605	2755	858184	31m 01s	33	11141	2673	784168	149	30m 27s
Illinois	4	8	0	144	0.08	4	8	0	144	0	0.08
Illinois_ca	3	3	1	48	0.02	3	3	1	48	0	0.02
Mesi	3	2	0	9	0.00	3	2	0	9	0	0.00
Mesi_ca	3	2	0	13	0.00	3	2	0	13	0	0.00
Moesi	3	2	0	10	0.01	3	2	0	10	0	0.01
Moesi_ca	3	2	0	13	0.00	3	2	0	13	0	0.00
Synapse	2	1	0	16	0.01	2	1	0	16	0	0.01
Xerox P.D.	7	13	0	388	0.23	7	13	0	388	0	0.23

performs better or outperforms (on the largest benchmarks) the model checkers PFS and Undip on the problems taken from their distributions. In addition, these two systems are not capable of handling many of the problems considered here such as those listed in the category “Imperative Programs” (their input syntax and the theoretical framework they are based on are too restrictive to accept them).

## 7. DISCUSSION

We have given a comprehensive account of our approach to the model checking of safety properties of infinite state systems manipulating array variables by SMT solving. The idea of using arrays to represent system states is not new in model-checking (see in particular [55, 54]); what seems to be new in our approach is the fully declarative characterization of *both* the topology and the (local) data structures of systems by using theories. This has two advantages. First, implementations of our approach can handle a wide range of topologies without modifying the underlying data structures representing sets of states. This is in contrast with recently developed techniques [2, 3] for the uniform verification of parametrized systems, which consist in exploring the state space of a system by using a

Table 3: Imperative Programs

Problem	Default Setting					Best Setting					
	d	#n	#d	#SMT	time	d	#n	#d	#SMT	#inv.	time
Find	4	27	7	691	0.90	4	27	7	691	0	0.90
Max_in_Array	-	-	-	-	timeout	2	1	1	46	5	0.03
Selection_Sort	-	-	-	-	timeout	5	13	2	1141	11	0.62
Strcat	-	-	-	-	timeout	2	2	2	80	2	0.07
Strcmp	-	-	-	-	timeout	2	1	1	21	3	0.01
Strcopy	3	3	1	694	1.22	3	3	2	564	4	0.38

Table 4: Miscellanea

Problem	Default Setting					Best Setting					
	d	#n	#d	#SMT	time	d	#n	#d	#SMT	#inv.	time
Alternating_bit	-	-	-	-	timeout	21	1008	156	41894	1	44.48
Bakery	6	12	0	86	0.04	6	12	0	86	0	0.04
Bakery2	6	22	1	247	0.07	6	22	1	247	0	0.07
Controller	6	8	0	95	0.03	6	8	0	95	0	0.03
Csm	-	-	-	-	timeout	2	2	2	76	1	0.02
Filter_simple	-	-	-	-	timeout	2	4	4	1013	132	3.94
Fischer	10	16	2	336	0.16	10	16	2	336	0	0.16
Fischer_U	8	13	3	198	0.08	8	13	3	198	0	0.08
German	26	2642	678	157870	191.39	26	2642	678	157870	0	191.39
Ins_sort	-	-	-	-	timeout	2	2	1	40	1	0.04
MIS	-	-	-	-	timeout	1	0	0	1261	95	0.85
Mux_Sem	7	15	0	174	0.04	7	15	0	174	0	0.04
Mux_Sem_param	4	5	0	85	0.04	2	3	1	57	4	0.02
Order	3	3	0	18	0.01	2	2	2	16	2	0.01
Simple	2	1	0	10	0.00	2	1	0	10	0	0.00
Swimming_Pool	3	81	0	1300	0.67	3	62	3	927	0	0.73
Szymanski+	21	685	102	43236	47.00	2	1	1	90	2	0.04
Ticket_o	-	-	-	-	timeout	3	4	2	201	10	0.06
Token_Ring	3	2	0	30	0.02	3	2	0	30	0	0.02
Tricky	8	7	0	22	0.02	2	1	1	13	1	0.00
Two_Semaphores	4	5	1	48	0.02	4	5	1	48	0	0.02

finitary representation of (infinite) sets of states and require substantial modifications in the computation of the pre-image to adapt to different topologies. Second, since SMT solvers are capable of handling several theories in combinations, we can avoid encoding everything in one theory, which has already been proved detrimental to performances in [19, 18]. SMT techniques were already employed in model-checking [24, 9], but only in the bounded case (whose aim is mostly limited at finding bugs, not at full verification).

In more details, our contributions are the following. First, we have explained how to use certain classes of first-order formulae to represent sets of states and identified the requirements to mechanize a fully symbolic and declarative version of backward reachability. Second, we have discussed sufficient conditions for the termination of the procedure on the theories used to specify the topology (indexes) and the data (elements) manipulated by the array-based system. Third, we have argued that the classes of formulae allow us to specify a variety of parametrized and distributed systems, and imperative algorithms manipulating arrays. Finally, we have studied invariant synthesis techniques and their integration in the backward reachability procedure. Theoretically, we have given sufficient conditions for the completeness on the theories of indexes and elements of the array-based system. Pragmatically, we have described how to interleave invariant guessing and backward reachability so as to ameliorate the termination of the latter. We have implemented the proposed techniques in MCMT and evaluated their viability on several benchmark problems extracted from a variety of sources. The experimental results have confirmed the efficiency and flexibility of our approach.

**7.1. Related work.** We now discuss the main differences and similarities with existing approaches to the verification of safety properties of infinite state systems. We believe it is convenient to recall two distinct and complementary approaches among the many possible alternatives available in the literature. In examining related works, we do not attempt to be exhaustive (we consider this an almost desperate task given the huge amount of work in this area) but rather to position our approach with respect to some of the main lines of research in the field.

The first approach is pioneered in [1] and its main notion is that of well-structured system. For example, it was implemented in two systems (see, e.g., [2, 3]), which were able to automatically verify several protocols for mutual exclusion and cache coherence. One of the key ingredients to the success of these tools is their capability to perform accurate fix-point checks so as to reduce the number of iterations of the backward search procedure. A fix-point check is implemented by ‘embedding’ an old configuration (i.e. a finite representation of a potentially infinite set of states) into a newly computed pre-image; if this is the case, then the new pre-image is considered “redundant” (i.e., not contributing new information about the set of backward reachable states) and thus can be discarded without loss of precision. Indeed, the exhaustive enumeration of embeddings has a high computational cost. An additional problem is that constraints are only used to represent the data manipulated by the system while its topology is encoded by *ad hoc* data structures. This requires to implement from scratch algorithms both to compute pre-images and embeddings, each time the topology of the systems to verify is modified. On the contrary, MCMT uses particular classes of *first-order formulae* to represent configurations parametrised with respect to a theory of the data and a theory of the topology of the system so that pre-image computation reduces to a fixed set of logical manipulations and fix-point checking to solve SMT problems containing universally quantified variables. To mechanize these tests, a quantifier-instantiation procedure is used, which is the logical counterpart of the enumeration of “embeddings.” Interestingly, this notion of “embedding” can be recaptured via classical model theory (see [37] or Section 4 above) in the logical framework underlying MCMT, a fact that allows us to import into our setting the decidability results of [1] for backward reachability. Another important advantage of the approach underlying MCMT over that proposed in [1] is its broader scope of applications with respect to the implementations in [2, 3, 4]. The use of theories for specifying the data and the topology allows one to model disparate classes of systems in a natural way. Furthermore, even if the quantifier instantiation procedure becomes incomplete with rich theories, it can soundly be used and may still permit to prove the safety of a system. In fact, MCMT has been successfully employed to verify sequential programs (such as sorting algorithms) that are far beyond the reach of the systems described in [2, 3].

The second and complementary approach to model checking infinite state system relies on *predicate abstraction* techniques, initially proposed in [44]. The idea is to abstract the system to one with finite states, to perform finite-state model checking, and to refine spurious traces (if any) by using decision procedures or SMT solvers. This technique has been implemented in several tools and is often combined with interpolation algorithms for the refinement phase. As pointed out in [34, 46], predicate abstraction must be carefully adapted when (universal) quantification is used to specify the transitions of the system or its properties, as it is the case for the problems tackled by MCMT. There are two crucial problems to be solved. The first is to find an appropriate set of predicates to compute the abstraction of the system. In fact, besides system variables, universally quantified variables may also

occur in the system. The second problem is that the computation of the abstraction as well as its refinement require to solve proof obligations containing universal quantifiers. Hence, we need to perform suitable quantifier instantiation in order to enable the use of decision procedures or SMT solving techniques for quantifier-free formulae. The first problem is solved by Skolemization [34] or fixing the number of variables in the system [46] so that standard predicate abstraction techniques can still be used. The second problem is solved by adopting very straightforward (sometimes naive) and incomplete quantifier instantiation procedures. While being computationally cheap and easy to implement, the heuristics used for quantifier instantiation are largely imprecise and does not permit the detection of redundancies due to variable permutations, internal symmetries, and so on. Experiments performed with MCMT, tuned to mimic these simple instantiation strategies, show much poorer performances. We believe that the reasons of success of the predicate abstraction techniques in [34, 46] lie in the clever heuristics used to find and refine the set of predicates for the abstraction. The current implementation of MCMT is orthogonal to the predicate abstraction approach; it features an extensive quantifier instantiation (which is complete for the theories over the indexes satisfying the Hypothesis (I) from Theorem 3.3 and is enhanced with completeness preserving heuristics to avoid useless instances), but it performs only a primitive form of predicate abstraction, called signature abstraction (see Section 6.3). Another big difference is how abstraction is used in MCMT: the set of backward reachable states is always computed precisely while abstraction is only exploited for guessing candidate invariants which are then used to prune the set of backward reachable states. Since we represent sets of states by formulae, guessing and then using the synthesized invariants turns out to be extremely easy, thereby helping to solve the tension between model checking and deductive techniques that has been discussed a lot in the literature and is still problematic in the tools described in [2, 3] where sets of states are represented by ad hoc data structures.

Besides the two main approaches mentioned above, there is a third line of research in the area that applied constraint solving techniques to the model-checking of infinite state systems. One of the first attempts was described in [19] and then furtherly studied in [18]. The idea was to use composite constraint domains (such as integers and Booleans) to encode the data and the control flow of, for example, instances of parametrised systems. Compared to our framework, the verification methods in [19, 18] are not capable of checking safety regardless of the number of process in a system but only supports the verification of its instances. Indeed, increasing the number of processes quickly degrades performances. Babylon is a tool for the verification of counting abstractions of parametrized systems (e.g., multithreaded Java programs [28]). It uses a graph-based data structure to encode disjunctive normal forms of integer arithmetic constraints. Computing pre-images requires computationally expensive normalization, which is not needed for us as SMT solvers efficiently handle arbitrary integer constraints. Brain is a model-checker for transition systems with finitely many integer variables which uses an incremental version of Hilbert's bases to efficiently perform entailment or satisfiability checking of integer constraints (the results reported in [56] shows that it scales very well). Taking  $T_I$  to be an enumerated data-type theory, the notion of array-based systems considered in this paper reduce to those used by Brain. However, many of the systems that can be modelled as array-based systems cannot be handled by Brain. Another interesting proposal to uniform verification of parametrized systems using constraint solving techniques is [15], where a decidability result for  $\Sigma_2^0$ -formulae is derived (these are  $\exists\forall$ -formulae roughly corresponding to those covered



by Theorem 3.3 above, for the special case in which the models of the theory  $T_I$  are all the finite linear orders). While the representation of states in [15] is (fully) declarative, transitions are not, as a rewriting semantics (with constraints) is employed. Since transitions are not declaratively handled, the task of proving pre-image closure becomes non trivial; in [15], pre-image closure of  $\Sigma_2^0$ -formulae under transitions encoded by  $\Sigma_2^0$ -formulae ensures the effectiveness of the tests for inductive invariant and bounded reachability analysis, but not for fix-point checks. In our approach, an easy (but orthogonal) pre-image closure result for existential state descriptions (under certain  $\Sigma_2^0$ -formulae representing transitions) gives the effectiveness of fix-point checks, thus allowing implementation of backward search.

**7.2. Future work.** We envisage to develop the work described here in three directions. First, we plan to enhance the implementation of the signature abstraction technique in future releases of MCMT. The idea is to find the best trade-off between the advantages of predicate abstraction and extensive quantifier instantiation. Another aspect is the design of methods for the dynamic refinement of the abstraction along the lines of the counter-example-guided-refinement (CEGAR) loop [44]. A complementary approach could be to use techniques for the automatic discovery of relationships among values of array elements developed in abstract interpretation (see, e.g., [43]). Second, we want to perform more extensive experiments for different classes of systems. For example, we have already started to investigate parametrised timed automata (introduced in [6]) with MCMT and found encouraging preliminary results [20]. Another class of problems in which successful experiments have been performed with MCMT concerns the verification of fault-tolerant distributed algorithms [8, 7]. The third line of future research consists of in exploring further and then implementing the verification method for a sub-class of liveness properties of array-based systems sketched in [37].

**Acknowledgments.** The authors would like to thanks the anonymous referees for the useful remarks that helped to improve the clarity of the paper.

The second author was partially supported by the FP7-ICT-2007-1 Project no. 216471, “AVANTSSAR: Automated Validation of Trust and Security of Service-oriented Architectures” ([www.avantssar.eu](http://www.avantssar.eu)) and by the Incoming Team Project “SIAM: Automated Security Analysis of Identity and Access Management Systems,” funded by Provincia Autonoma di Trento in the context of the COFUND action of the European Commission (FP7).

## REFERENCES

- [1] P. A. Abdulla, K. Cerans, B. Jonsson, and Y.-K. Tsay. General decidability theorems for infinite-state systems. In *Proc. of LICS*, pages 313–321, 1996.
- [2] P. A. Abdulla, G. Delzanno, N. B. Henda, and A. Rezine. Regular model checking without transducers. In *TACAS*, volume 4424 of *LNCS*, pages 721–736, 2007.
- [3] P. A. Abdulla, G. Delzanno, and A. Rezine. Parameterized verification of infinite-state processes with global conditions. In *CAV*, volume 4590 of *LNCS*, pages 145–157, 2007.
- [4] Parosh Aziz Abdulla, Noomene Ben Henda, Giorgio Delzanno, and Ahmed Rezine. Handling parameterized systems with non-atomic global conditions. In *Proc. of VMCAI*, volume 4905 of *LNCS*, pages 22–36, 2008.
- [5] Parosh Aziz Abdulla and Bengt Jonsson. Verifying programs with unreliable channels. *Information and Computation*, 127(2):91–101, 1996.
- [6] Parosh Aziz Abdulla and Bengt Jonsson. Model checking of systems with many identical timed processes. *Theoretical Computer Science*, pages 241–264, 2003.

- [7] F. Alberti, S. Ghilardi, E. Pagani, S. Ranise, and G. P. Rossi. Automated Support for the Design and Validation of Fault Tolerant Parameterized Systems: a case study. In *Proc. of AVOCS 10*, Electr. Comm. of the EASST, 2010.
- [8] F. Alberti, S. Ghilardi, E. Pagani, S. Ranise, and G. P. Rossi. Brief Announcement: Automated Support for the Design and Validation of Fault Tolerant Parameterized Systems—a case study. In *Proc. of DISC 10*, number 6343 in LNCS, pages 392–394, 2010.
- [9] A. Armando, J. Mantovani, and L. Platania. Bounded Model Checking of Software using SMT Solvers instead of SAT Solvers. In *Proc. of SPIN’06*, number 3925 in LNCS, pages 146–162, 2006.
- [10] F. Baader and T. Nipkow. *Term Rewriting and All That*. Cambridge University Press, United Kingdom, 1998.
- [11] Franz Baader and Silvio Ghilardi. Connecting many-sorted theories. *Journal of Symbolic Logic*, 72:535–583, 2007.
- [12] Ittai Balaban, Yi Fang, Amir Pnueli, and Lenore Zuck. Iiv: An invisible invariant verifier. In *Computer Aided Verification (CAV) 2005*, 2005.
- [13] D. Beyer, T. A. Henzinger, R. Majumdar, and A. Rybalchenko. Invariant Synthesis for Combined Theories. In *VMCAI’07*, volume 4349 of LNCS, 2007.
- [14] N. Bjørner, A. Browne, and Z. Manna. Automatic Generation of Invariants and Assertions. In *Principles and Practice of Constraint Programming - CP’95, First International Conference, CP’95, Cassis, France*, volume 976 of LNCS, pages 589–623. Springer, 1995.
- [15] A. Bouajjani, P. Habermehl, Y. Yurski, and M. Sighireanu. Rewriting systems with data. In *Proc. of Symp. on Fund. of Comp. Th. (FCT 07)*, pages 1–22, 2007.
- [16] Aaron R. Bradley and Zohar Manna. Property-Directed Incremental Invariant Generation. *Formal Aspects of Computing*, 2009. To appear.
- [17] Aaron R. Bradley, Zohar Manna, and Henny B. Sipma. What’s decidable about arrays? In *Proc. of VMCAI*, volume 3855 of LNCS, pages 427–442, 2006.
- [18] T. Bultan, R. Gerber, and C. League. Composite model-checking: verification with type-specific symbolic representations. *ACM Trans. on Soft. Eng. and Meth.*, 9(1):3–50, 2000.
- [19] T. Bultan, R. Gerber, and W. Pugh. Model-checking concurrent systems with unbounded integer variables: symbolic representations, approximations, and experimental results. *ACM Trans. on Progr. Lang. and Sys.*, 21(4):747–789, 1999.
- [20] A. Carioni, S. Ghilardi, and S. Ranise. MCMT in the Land of Parametrized Timed Automata. In *Proc. of VERIFY 10*, 2010.
- [21] A. Chagrov and M. Zakharyashev. *Modal Logic*. Clarendon Press, 1997.
- [22] Chen-Chung Chang and Jerome H. Keisler. *Model Theory*. North-Holland, Amsterdam-London, third edition, 1990.
- [23] L. de Moura and N. Bjørner. Efficient e-matching for smt solvers. In *Proc. of CADE*, LNCS, 2007.
- [24] L. de Moura, H. Rueß, and M. Sorea. Lazy theorem proving for bounded model checking over infinite domains. In *Proc. CADE*, volume 2392 of LNCS, 2002.
- [25] D. Déharbe and S. Ranise. Satisfiability solving for software verification. *Int. Journal on STTT*, volume 11, number 3, 2009.
- [26] G. Delzanno. Automatic verification of parameterized cache coherence protocols. In *Proc. of CAV*, number 1855 in LNCS, 2000.
- [27] G. Delzanno, J. Esparza, and A. Podelski. Constraint-based analysis of broadcast protocols. In *Proc. of CSL*, volume 1683 of LNCS, pages 50–66, 1999.
- [28] G. Delzanno, J.-F. Raskin, and L. Van Begin. Towards the automated verification of multi-threaded java programs. In *8th Int. Conf. on TACAS*, number 2280 in LNCS, 2002.
- [29] E. W. Dijkstra. Guarded commands, nondeterminacy and formal derivation of programs. *Communications of the ACM*, 18:453–457, 1975.
- [30] D. L. Dill and H. Wong-Toi. Verification of Real-Time Systems by Successive Over and Under Approximation. In *Computer Aided Verification, 7th International Conference, Liège, Belgium*, volume 939 of LNCS, pages 409–422. Springer, 1995.
- [31] Bruno Dutertre and Leonardo De Moura. The yices smt solver. Technical report, Computer Science Laboratory, SRI International, 2006. Available at <http://yices.csl.sri.com>.
- [32] Herbert B. Enderton. *A Mathematical Introduction to Logic*. Academic Press, New York-London, 1972.

- [33] J. Esparza, A. Finkel, and R. Mayr. On the verification of broadcast protocols. In *Proc. of LICS*, pages 352–359. IEEE Computer Society, 1999.
- [34] C. Flanagan and S. Qadeer. Predicate abstraction for software verification. In *Proc. of POPL'02*, pages 191–202. ACM, 2002.
- [35] J. Gallier. What's so Special about Kruskal's Theorem and the Ordinal  $\Gamma_0$ ? A Survey of Some Results in Proof Theory. *Annals of Pure and Applied Logic*, 53:199–260, 1991.
- [36] Y. Ge, C. Barrett, and C. Tinelli. Solving quantified verification conditions using satisfiability modulo theories. In *Proc. of CADE-21*, LNCS, 2007.
- [37] S. Ghilardi, E. Nicolini, S. Ranise, and D. Zucchelli. Towards SMT Model-Checking of Array-based Systems. In *Proc. of IJCAR*, LNCS, 2008. Extended version available online as Tech. Report RI318-08 at <http://homes.dsi.unimi.it/~zucchelli/publications/techreport/GhiNiRaZu-RI318-08.pdf>.
- [38] S. Ghilardi and S. Ranise. A Note on the Stopping Failures Models. 2009. Unpublished Draft, available from MCMT web site.
- [39] S. Ghilardi and S. Ranise. Goal Directed Invariant Synthesis for Model Checking Modulo Theories. In *(TABLEAUX 09)*, LNAI, pages 173–188. Springer, 2009.
- [40] S. Ghilardi and S. Ranise. Model Checking Modulo Theory at work: the integration of Yices in MCMT. In *AFM 09 (co-located with CAV09)*, 2009.
- [41] S. Ghilardi and S. Ranise. MCMT: a Model Checker Modulo Theories. In *Proc. of IJCAR'10*, LNCS, 2010. To appear.
- [42] S. Ghilardi, S. Ranise, and T. Valsecchi. Light-Weight SMT-based Model-Checking. In *Proc. of AVOCS 07-08*, ENTCS, 2008.
- [43] D. Gopan, T. Reps, and M. Sagiv. Numeric analysis of array operations. In *Conference Record of the Thirty-Second ACM Symposium on Principles of Programming Languages, (Long Beach, CA)*, 338-350, 2005.
- [44] S. Graf and H. Saïdi. Construction of abstract state graphs with PVS. In *Proc. of CAV 1997*, volume 1254 of LNCS. Springer, 1997.
- [45] T. A. Henzinger, P.-H. Ho, and H. Wong-Toi. HYTECH: A Model Checker for Hybrid Systems. In *Computer Aided Verification, 9th International Conference, CAV '97, Haifa, Israel*, volume 1254 of LNCS, pages 460–463. Springer, 1997.
- [46] S. K. Lahiri and R. E. Bryant. Predicate abstraction with indexed predicates. *ACM Transactions on Computational Logic (TOCL)*, 9(1), 2007.
- [47] Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.
- [48] A. Rezne P. A. Abdulla, G. Delzanno. Approximated context-sensitive analysis for parametrized verification. In *Proc. of FORTE 09*, LNCS, 2009.
- [49] S. Park and D. L. Dill. Verification of FLASH Cache Coherence Protocol by Aggregation of Distributed Transactions. In *In Proceedings of the 8th Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 288–296. ACM Press, 1996.
- [50] Schnoebelen Philippe. Verifying lossy channel systems has nonprimitive recursive complexity. *Information Processing Letters*, 83(5):251–261, 2002.
- [51] A. Pnueli, S. Ruath, and L. D. Zuck. Automatic deductive verification with invisible invariants. In *Proc. of TACAS 2001*, volume 2031 of LNCS, 2001.
- [52] S. Ranise and C. Tinelli. The SMT-LIB Standard: Version 1.2. Technical report, Dep. of Comp. Science, Iowa, 2006. Available at <http://www.SMT-LIB.org/papers>.
- [53] Silvio Ranise and Cesare Tinelli. The SMT-LIB standard: Version 1.2. Technical report, 2006. Available at <http://combination.cs.uiowa.edu/smtlib/papers/format-v1.2-r06.08.30.pdf>.
- [54] A. W. Roscoe, R. S. Lazic, and T. C. Newcomb. On model checking data-independent systems with arrays without reset. *Theory and Practice of Logic Programming*, pages 659–693, 2004.
- [55] A. W. Roscoe, R. S. Lazic, and Tom Newcomb. On model checking data-independent systems with arrays with whole-array operations. In *Communicating Sequential Processes*. Springer LNCS, 2005.
- [56] T. Rybina and A. Voronkov. Using canonical representations of solutions to speed up infinite-state model checking. In *Proc. of CAV*, number 2404 in LNCS, 2002.

## APPENDIX A. OMITTED PROOFS

**Decidability of restricted satisfiability checking.** The following result is a simple generalization of Theorem 3.3 (of Section 3.2).

**Theorem A.1.** *The  $A_I^E$ -satisfiability of a sentence of the kind*

$$\exists a_1 \cdots \exists a_n \exists \underline{i} \exists \underline{e} \forall \underline{j} \psi(\underline{i}, \underline{j}, \underline{e}, a_1[\underline{i}], \dots, a_n[\underline{i}], a_1[\underline{j}], \dots, a_n[\underline{j}]) \quad (\text{A.1})$$

*is decidable. Moreover, the following conditions are equivalent:*

- (i) *the sentence (A.1) is  $A_I^E$ -satisfiable;*
- (ii) *the sentence (A.1) is satisfiable in a finite index model of  $A_I^E$ ;*
- (iii) *the sentence*

$$\exists a_1 \cdots \exists a_n \exists \underline{i} \exists \underline{e} \bigwedge_{\sigma} \psi(\underline{i}, \underline{j}\sigma, \underline{e}, a_1[\underline{i}], \dots, a_n[\underline{i}], a_1[\underline{j}\sigma], \dots, a_n[\underline{j}\sigma]) \quad (\text{A.2})$$

*is  $A_I^E$ -satisfiable (here  $\sigma$  ranges onto the substitutions mapping the variables  $\underline{j}$  into the set of representative  $\Sigma_I(\underline{i})$ -terms).*

*Proof.* In order to avoid difficulties with the notation, we consider the case where  $n = 1$  only (the reader may check that there is no loss of generality in that).<sup>17</sup> We first show that the  $A_I^E$ -satisfiability of

$$\exists a \exists \underline{i} \exists \underline{e} \forall \underline{j} \psi(\underline{i}, \underline{j}, \underline{e}, a[\underline{i}], a[\underline{j}]) \quad (\text{A.3})$$

is equivalent to the  $A_I^E$ -satisfiability of

$$\exists a \exists \underline{i} \exists \underline{e} \bigwedge_{\sigma} \psi(\underline{i}, \underline{j}\sigma, \underline{e}, a[\underline{i}], a[\underline{j}\sigma]) \quad (\text{A.4})$$

where  $\sigma$  ranges onto the substitutions mapping the variables  $\underline{j}$  into the set of representative  $\Sigma_I(\underline{i})$ -terms.

That  $A_I^E$ -satisfiability of (A.3) implies  $A_I^E$ -satisfiability of (A.4) follows from trivial logical manipulations, so let's assume  $A_I^E$ -satisfiability of (A.4) and show  $A_I^E$ -satisfiability of (A.3). Let  $\mathcal{M}$  be a model of (A.4); we can assign elements in this model to the variables  $a, \underline{i}, \underline{e}$  in such a way that (under such an assignment  $\mathbf{a}$ ) we have  $\mathcal{M}, \mathbf{a} \models \bigwedge_{\sigma} \psi(\underline{i}, \underline{j}\sigma, \underline{e}, a[\underline{i}], a[\underline{j}\sigma])$ . Consider the model  $\mathcal{N}$  which is obtained from  $\mathcal{M}$  by restricting the interpretation of the sort **INDEX** (and of all function and relation symbols for indexes) to the  $\Sigma_I$ -substructure generated by the elements assigned by  $\mathbf{a}$  to the  $\underline{i}$ : since models of  $T_I$  are closed under substructures, this substructure is a model of  $T_I$  and consequently  $\mathcal{N}$  is still a model of  $A_I^E$ . Now let  $s$  be the restriction of  $\mathbf{a}(a)$  to the new smaller index domain and let  $\tilde{\mathbf{a}}$  be the assignment differing from  $\mathbf{a}$  only for assigning  $s$  to  $a$  (instead of  $\mathbf{a}(a)$ ); since  $\psi$  is quantifier free and since, varying  $\sigma$ , the elements assigned to the terms  $\underline{j}\sigma$  covers all possible  $\underline{j}$ -tuples of elements in the interpretation of the sort **INDEX** in  $\mathcal{N}$ , we have  $\mathcal{N}, \tilde{\mathbf{a}} \models \forall \underline{j} \psi(\underline{i}, \underline{j}, \underline{e}, a[\underline{i}], a[\underline{j}])$ . This shows that  $\mathcal{N} \models \exists a \exists \underline{i} \forall \underline{j} \psi(\underline{i}, \underline{j}, a[\underline{i}], a[\underline{j}])$ , i.e that (A.3) holds. Notice that  $\mathcal{N}$  is a finite index model,<sup>18</sup> hence we proved also the equivalence between (i) and (ii).

<sup>17</sup>Since existentially quantifying over variables that do not occur in the formula does not affect satisfiability, we can also assume that the tuple  $\underline{i}$  is not empty (this observation is needed if we want to prevent the structure  $\mathcal{N}$  defined below from having empty index domain).

<sup>18</sup>This is because  $T_I$  is locally finite and the  $\Sigma_I$  reduct of  $\mathcal{N}$  is a structure which is generated by finitely many elements.

We now need to decide  $A_I^E$ -satisfiability of sentences (A.4). Let  $\underline{t}$  be the representative  $\Sigma(\underline{i})$ -terms and let us put them in bijective correspondence with fresh variables  $\underline{l}$  of sort INDEX; let  $\psi_\sigma(\underline{i}, \underline{l}, \underline{e}, a[\underline{i}], a[\underline{l}])$  be the formula obtained by replacing in  $\psi(\underline{i}, \underline{j}\sigma, \underline{e}, a[\underline{i}], a[\underline{j}\sigma])$  the  $\Sigma(\underline{i})$ -terms  $\underline{j}\sigma$  by the corresponding  $\underline{l}$ . We first rewrite (A.4) as

$$\exists a \exists \underline{i} \exists \underline{e} \exists \underline{l} (\underline{l} = \underline{t} \wedge \bigwedge_{\sigma} \psi_\sigma(\underline{i}, \underline{l}, \underline{e}, a[\underline{i}], a[\underline{l}])) \quad (\text{A.5})$$

(here  $\underline{l} = \underline{t}$  means component-wise equality, expressed as a conjunction).

Notice that  $T_I$  and  $T_E$  are disjoint (they do not have even any sort in common), which means that  $\underline{l} = \underline{t} \wedge \bigwedge_{\sigma} \psi_\sigma(\underline{i}, \underline{l}, \underline{e}, a[\underline{i}], a[\underline{l}])$  is a Boolean combination of  $\Sigma_I$ -atoms and of  $\Sigma_E$ -atoms (in the latter kind of atoms, the variables for elements are replaced by the terms  $a[\underline{i}], a[\underline{l}]$ ). This means that our decision problem can be further rephrased in terms of the problem of deciding for  $A_I^E$ -satisfiability formulae like

$$\psi_I(\underline{j}) \wedge a[\underline{j}] = \underline{d} \wedge \psi_E(\underline{d}, \underline{e}) \quad (\text{A.6})$$

where  $\psi_I(\underline{j})$  is a conjunction of  $\Sigma_I$ -literals and  $\psi_E(\underline{d}, \underline{e})$  is a conjunction of  $\Sigma_E$ -literals.

Since we are looking for a model of  $T_I$ , a model of  $T_E$  and for a function connecting their domains (the function interpreting the variable  $a$ ), this is a satisfiability problem for a theory connection (in the sense of [11]):<sup>19</sup> since the signatures of  $T_I, T_E$  are disjoint, the problem is decided by propagating equalities.<sup>20</sup> Hence, to decide (A.6), it is sufficient to apply the following steps:

- guess an equivalence relation  $\Pi$  on the index variables  $\underline{j}$  (let's assume  $\underline{j} = j_1, \dots, j_n$ );
- check  $\psi_I(\underline{j}) \cup \{j_k = j_l \mid (j_k, j_l) \in \Pi\} \cup \{j_k \neq j_l \mid (j_k, j_l) \notin \Pi\}$  for  $T_I$ -satisfiability;
- check  $\psi_E(\underline{d}, \underline{e}) \cup \{d_k = d_l \mid (j_k, j_l) \in \Pi\}$  for  $T_E$ -satisfiability;
- return ‘unsatisfiable’ iff failure is reported in the previous two steps for all possible guesses.

Soundness and completeness of the above procedure are easy. □

**Undecidability of backward reachability.** Here, we give the *proof of Theorem 4.1* (of Section 4.2).

*Proof.* A *two registers Minsky machine* is a finite set  $\mathbf{P}$  of instructions (also called a program) for manipulating configurations seen as triples  $(q, m, n)$  where  $m, n$  are natural numbers representing the registers content and  $q$  represents the machine location state ( $q$  varies on a fixed finite set  $Q$ ). There are four possible kinds of instructions, inducing transformations on the configurations as explained in Table 5. A  $\mathbf{P}$ -transformation is a transformation induced by an instruction of  $\mathbf{P}$  on a certain configuration. For a Minsky machine  $\mathbf{P}$ , we write  $(q, m, n) \rightarrow_{\mathbf{P}}^* (q', m', n')$  to say that it is possible to reach configuration  $(q', m', n')$  from  $(q, m, n)$  by applying finitely many  $\mathbf{P}$ -transformations. Given a Minsky machine  $\mathbf{P}$  and an initial configuration  $(q_0, m_0, n_0)$ , the problem of checking whether a configuration  $(q', m', n')$  is reachable from  $(q_0, m_0, n_0)$  (i.e., if  $(q_0, m_0, n_0) \rightarrow_{\mathbf{P}}^* (q', m', n')$  holds or not) is

<sup>19</sup>Strictly speaking, one cannot directly apply the results from [11], because in this paper we have adopted a ‘semantic’ notion of theory characterized by a class  $\mathcal{C}$  of models. In other words, the class  $\mathcal{C}$  of models is not required to be elementary.

<sup>20</sup>This is different from the standard Nelson-Oppen combination, where also inequalities must be propagated.

N.	Instruction	Transformation
I	$q \rightarrow (r, 1, 0)$	$(q, m, n) \rightarrow (r, m + 1, n)$
II	$q \rightarrow (r, 0, 1)$	$(q, m, n) \rightarrow (r, m, n + 1)$
III	$q \rightarrow (r, -1, 0)[r']$	<b>if</b> $m \neq 0$ <b>then</b> $(q, m, n) \rightarrow (r, m - 1, n)$ <b>else</b> $(q, m, n) \rightarrow (r', m, n)$
IV	$q \rightarrow (r, 0, -1)[r']$	<b>if</b> $n \neq 0$ <b>then</b> $(q, m, n) \rightarrow (r, m, n - 1)$ <b>else</b> $(q, m, n) \rightarrow (r', m, n)$

Table 5: Instructions and related transformations for (two-registers) Minsky Machines

called *the (second) reachability (configuration) problem*. It is well-known<sup>21</sup> that there exists a (two-register) Minsky machine  $\mathbf{P}$  and a configuration  $(q_0, m_0, n_0)$  such that the second reachability configuration problem is undecidable. To simplify the matter, we assume that  $m_0 = 0$  and  $n_0 = 0$ : there is no loss of generality in that, because one can add to the program  $\mathbf{P}$  more states and instructions (precisely  $m_0 + n_0$  further states and instructions of type I-II) for the initialization to  $m_0, n_0$ .

We build a locally finite array-based system  $\mathcal{S}_{\mathbf{P}} = (a, I_{\mathbf{P}}, \tau_{\mathbf{P}})$  and an  $\exists^I$ -formula  $U_{q,m,n}$  such that  $\mathcal{S}$  is unsafe w.r.t.  $U_{q,m,n}$  iff the machine  $\mathbf{P}$  reaches the configuration  $(q, m, n)$ . We take as  $\Sigma_I$  the signature having two constants  $o, o'$  and a binary relation  $S$ ; models of  $T_I$  are the  $\Sigma_I$ -structures satisfying the axioms

$$\begin{aligned} \forall i \neg S(i, o), \quad \forall i \forall j_1 \forall j_2 (S(i, j_1) \wedge S(i, j_2) \rightarrow j_1 = j_2), \\ S(o, o'), \quad \forall i_1 \forall i_2 \forall j (S(i_1, j) \wedge S(i_2, j) \rightarrow i_1 = i_2), \end{aligned}$$

saying that  $S$  is an injective partial function having  $o$  in the domain but not in the range. As  $\Sigma_E$  we take the enumerated datatype theory relative to the finite set  $Q \times \{0, 1\} \times \{0, 1\}$ . Notice that  $T_I, T_E$  are both locally finite; in addition,  $T_I$  is closed under substructures and  $T_E$  has quantifier elimination.

The idea is that of encoding a configuration  $(q, m, n)$  as any configuration  $s$  (in the formal sense of Subsection 4.1) satisfying the following conditions:

- (i) the support of  $s_I$  contains a substructure of the kind

$$o = i_0 \rightarrow_S o' = i_1 \rightarrow_S i_2 \cdots \rightarrow_S i_K$$

for some  $K > m, n$  (we write  $i \rightarrow_S j$  to mean that  $(i, j)$  is in the interpretation of the relational symbol  $S$  in  $s_I$ ).

- (ii) for all  $i$  in the support of  $s_I$ , if  $s(i) = \langle r, u, v \rangle$  then (a)  $r = q$ ; (b)  $u = 1$  iff  $i = i_k$  for a  $k \leq m$ ; (c)  $v = 1$  iff  $i = i_k$  for a  $k \leq n$ .

In case the above conditions (i)-(ii) hold, we say that  $s$  *bi-simulates*  $(q, m, n)$ .

The initial formula  $I$  is

$$\forall i ((i \neq o \wedge a[i] = \langle q_0, 0, 0 \rangle) \vee (i = o \wedge a[i] = \langle q_0, 1, 1 \rangle)).$$

Clearly for every model  $\mathcal{M}$  and for every  $s \in \text{ARRAY}^{\mathcal{M}}$ , the following happens:

- ( $\alpha$ ):  $\mathcal{M} \models I(s)$  iff  $s$  bi-simulates the initial machine configuration  $(q_0, 0, 0)$ .

We write the transition  $\tau$  in such a way that for every model  $\mathcal{M}$  and for every  $s, s' \in \text{ARRAY}^{\mathcal{M}}$ , the following happens:

<sup>21</sup>For details and further references, see for instance [21].

( $\beta$ ): if  $s$  bi-simulates  $(q, m, n)$ , then  $\mathcal{M} \models \tau(s, s')$  iff there is  $(q', m', n')$  such that  $s'$  bi-simulates  $(q', m', n')$  and  $(q, m, n) \rightarrow_{\mathbf{P}} (q', m', n')$ .

This goal is obtained by taking  $\tau$  to be a disjunction of  $T$ -formulae corresponding to the instructions for  $\mathbf{P}$ . The  $T$ -formula corresponding to the first kind of instructions  $q \rightarrow (r, 1, 0)$  is the following:<sup>22</sup>

$$\begin{aligned} \exists i_1 \exists i_2 \exists i_3 \quad & (S(i_1, i_2) \wedge S(i_2, i_3) \wedge pr_1(a[i_1]) = q \wedge pr_2(a[i_1]) = 1 \wedge \\ & \wedge pr_2(a[i_2]) = 0 \wedge pr_2(a[i_3]) = 0 \wedge a' = \lambda j F) \end{aligned}$$

where

$$\begin{aligned} F \quad := \quad & \text{if } (j = i_2) \text{ then } \langle r, 1, pr_3(a[j]) \rangle \\ & \text{else } \langle r, pr_2(a[j]), pr_3(a[j]) \rangle \end{aligned}$$

Instructions  $q \rightarrow (r, -1, 0)[r']$  of the kind (III) are simulated by the following  $T$ -formula

$$\exists i_1 \exists i_2 (S(i_1, i_2) \wedge pr_1(a[i_1]) = q \wedge pr_2(a[i_1]) = 1 \wedge pr_2(a[i_2]) = 0)$$

where

$$\begin{aligned} F \quad := \quad & \text{if } (i_1 \neq o \wedge j = i_1) \text{ then } \langle r, 0, pr_3(a[j]) \rangle \\ & \text{else if } (i_1 \neq o \wedge j \neq i_1) \text{ then } \langle r, pr_2(a[j]), pr_3(a[j]) \rangle \\ & \text{else } \langle r', pr_2(a[j]), pr_3(a[j]) \rangle \end{aligned}$$

$T$ -formulae for instructions of kind (II) and (IV) are defined accordingly.

We write the unsafe states formula  $U_{q,m,n}$  in such a way that for every model  $\mathcal{M}$  and for every  $s \in \text{ARRAY}^{\mathcal{M}}$ , the following happens:

( $\gamma$ ): if  $\mathcal{M} \models U_{q,m,n}(s)$  and  $s$  bi-simulates some machine configuration, then it bi-simulates  $(q, m, n)$ .

This goal is achieved by taking  $U_{q,m,n}$  to be the following formula (suppose  $m \geq n$ , the case  $n \leq m$  is symmetric):

$$\begin{aligned} \exists i_0 \cdots \exists i_{m+1} \quad & (i_0 = o \wedge \bigwedge_{0 \leq k \leq m} S(i_k, i_{k+1}) \wedge \bigwedge_{0 \leq k \leq n} a[i_k] = \langle q, 1, 1 \rangle \wedge \\ & \wedge \bigwedge_{n < k \leq m} a[i_k] = \langle q, 1, 0 \rangle \wedge a[i_{m+1}] = \langle q, 0, 0 \rangle). \end{aligned}$$

From ( $\alpha$ )-( $\beta$ )-( $\gamma$ ) above it is clear that  $\mathbf{P}$  reaches the configuration  $(q, m, n)$  iff  $\mathcal{S}$  is unsafe w.r.t.  $U_{q,m,n}$ , so that the latter is not decidable (for the left to right implication, take a run in a model with a large enough  $S$ -chain starting with  $o$ ).  $\square$

**Undecidability of unrestricted satisfiability checking.** We show that *Hypothesis (I) cannot be removed from the statement of Theorem 3.3* (and of Theorem A.1). We use a reduction to the reachability problem for Minsky machines as we have done for the proof of the undecidability of the safety problem (Theorem 4.1); the argument is similar to one used in [17].

Let  $T_I$  be the theory having as a class of models the natural numbers in the signature with just zero, the successor function, and  $\leq$ . Notice that this is not locally finite. Let  $T_E$  be the theory having  $Q \times \mathbb{N} \times \mathbb{N}$  as a unique structure. Here  $Q$  is like in the previous subsection of this Appendix. In the following, we freely use projections, sums, numerals, subtraction,

<sup>22</sup>For simplicity, we assume that the signature  $\Sigma_E$  is 4-sorted and endowed with the three projection functions  $pr_1, pr_2, pr_3$  mapping a data  $\langle r, u, v \rangle$  to  $r, u, v$ , respectively: there is no need of this assumption, but without it specifications become cumbersome.

constants for elements of  $Q$ , etc. Formally, all these operations can be defined in many ways and the precise way is not relevant for the argument below. In other words, we can avoid to define precisely the signature  $\Sigma_E$ . This sloppiness is justified because we must use a  $T_I$  not satisfying the local finiteness requirement from Hypothesis (I) of Theorem 3.3, whereas we can use an arbitrary  $T_E$ . Let  $\tau(a[j_1], a[j_2])$  abbreviate the disjunction of the following formulae describing the transformations from Table 5:

$$\begin{aligned} pr_2(a[j_1]) = q \wedge a[j_2] = \langle r, pr_2(a[j_1]) + 1, pr_3(a[j_1]) \rangle \\ pr_2(a[j_1]) = q \wedge a[j_2] = \langle r, pr_2(a[j_1]), pr_3(a[j_1]) + 1 \rangle \\ pr_2(a[j_1]) = q \wedge pr_2(a[j_1]) > 0 \wedge a[j_2] = \langle r, pr_2(a[j_1]) - 1, pr_3(a[j_1]) \rangle \\ pr_2(a[j_1]) = q \wedge pr_2(a[j_1]) = 0 \wedge a[j_2] = \langle r', pr_2(a[j_1]), pr_3(a[j_1]) \rangle \\ pr_2(a[j_1]) = q \wedge pr_3(a[j_1]) > 0 \wedge a[j_2] = \langle r, pr_2(a[j_1]), pr_3(a[j_1]) - 1 \rangle \\ pr_2(a[j_1]) = q \wedge pr_3(a[j_1]) = 0 \wedge a[j_2] = \langle r', pr_2(a[j_1]), pr_3(a[j_1]) \rangle \end{aligned}$$

Now consider the satisfiability of the following  $\exists^A, \forall^I$ -formula:

$$\begin{aligned} \exists a \exists i \exists j (i = 0 \wedge a[i] = \langle q_0, 0, 0 \rangle \wedge a[j] = \langle q, m, n \rangle \wedge \\ \wedge \forall j_1 \forall j_2 (j_1 < j \wedge j_2 = j_1 + 1 \rightarrow \tau(a[j_1], a[j_2]))) \end{aligned}$$

Clearly, this is satisfiable iff the configuration  $\langle q, m, n \rangle$  is reachable: the array  $a$  in fact stores the whole computation leading to  $\langle q, m, n \rangle$ . Thus satisfiability of  $\exists^A, \forall^I$ -formulae can be undecidable if  $T_I$  is not locally finite, even if the  $SMT(T_I)$  and the  $SMT(T_E)$  problems are decidable (and even if  $T_I$  is closed under substructures).

A final observation is crucial. If we keep local finiteness and drop closure under substructures in the statement of Hypothesis (I) from Theorem 3.3, then the above counterexample still applies! In fact, the successor function for indexes is used only in  $j_2 = j_1 + 1$  occurring in the formula above: we can replace the application of the successor function with a binary relation  $S(j_1, j_2)$  so as to recover local finiteness. However, closure under substructures is dropped as the structure of natural numbers has proper substructures if successor is a relation and not a function and these substructures must be excluded for the above argument to work (from satisfiability in such substructures a full computation cannot be recovered). Thus, we can conclude that *the two conditions of Hypothesis (I) are strictly connected and both needed.*